



FINAL REPORT

Development and Dissemination of a High-Resolution National Climate Change Dataset

Katharine Hayhoe, Texas Tech University (PI)

With contributions from Anne Stoner, Xiaohui Yang, Caleb Crow, Ranjini Swaminathan, Ian Scott-Fleming, Jung-Hee Ryu, Rodica Gelca, Sharmistha Swain, and the Texas Tech University High Performance Computing Center

Agreement Number G10AC00582

Date of Report: March 22, 2013

Period of Time Covered by Report: August 1, 2009 to March 31, 2013

Actual total cost: \$149,368.00

PUBLIC SUMMARY

Climate change is a global problem whose impacts are experienced at the local to regional scale. For this reason, the first step in assessing the impacts of climate change—on a species or an ecosystem, on water resources, or on an aspect of human society such as energy demand or agriculture—is often to develop projections of how temperature, precipitation, and other important aspects of climate might be expected to change in the future at the location of interest.

Global climate models produce future climate projections that are usually too coarse to capture the local characteristics that determine climate at any given location: the proximity of that location to a large body of water, for example, which would moderate extreme temperatures; or whether the location is in the arid “rain shadow” of a mountain.

In this project, we used an advanced statistical downscaling method that combines high-resolution observations with outputs from 16 different global climate models based on 4 future emission scenarios to generate the most comprehensive dataset of daily temperature and precipitation projections available for climate change impacts in the U.S. The gridded dataset covers the continental United States, southern Canada and northern Mexico at one-eighth degree resolution and Alaska at one-half degree resolution. We also quality-controlled observations from over 10,000 long-term weather stations and generated projections for each of these locations.

The high-resolution projections produced by this work have been rigorously quality-controlled for both errors and biases in the global climate and statistical downscaling models. We also calculated projected future changes in a broad range of impact-relevant indicators, from seasonal temperature to extreme precipitation days. The results of the error and bias tests and the indicator calculations are made available as part of this database.

TABLE OF CONTENTS

TECHNICAL SUMMARY	4
PURPOSE AND OBJECTIVES	6
ORGANIZATION AND APPROACH	8
Historical and Future Climate Scenarios	9
Global Climate Model Simulations	12
Downscaling Model	14
Gridded Observations	20
Station Observations	21
Computational Methods	22
PROJECT RESULTS	22
ANALYSIS AND FINDINGS	30
Improvements to downscaling inputs.....	30
Improvements to downscaling methodology	42
Improvements to computational methods	52
Error and bias analysis	53
Educating the user and the downscaler	59
CONCLUSIONS AND RECOMMENDATIONS.....	62
OUTREACH	63
REFERENCES	64
APPENDIX A: Downscaled Weather Stations: Latitude, Longitude and Station ID	
APPENDIX B: Regional Temperature-Precipitation Projections: Scatter Plots	
APPENDIX C: U.S. Climate Indicators: Global Mean Temperature Change Maps	
APPENDIX D: Evaluating Downscaled Projections: Error Maps	
APPENDIX E: Evaluating Downscaled Projections: Bias Maps	
APPENDIX F: ARRM Downscaling Method Documentation: Stoner et al. 2012	

TECHNICAL SUMMARY

This project has produced a **database of high-resolution projections of daily maximum and minimum temperature and 24-hour cumulative precipitation** suitable for use as input to a range of ecological, hydrological, and societal studies to assess the impacts of climate change at the local to regional scale. The database consists of daily projections from 1960 to 2099 from 16 global models archived by the Coupled Model Intercomparison Project version 3 (CMIP3) and will be supplemented by new model simulations from CMIP5 as available. Future CMIP3 projections are based on four different emission scenarios (A1fi higher, A2 mid-high, A1B mid-low and B1 lower) from the Intergovernmental Panel on Climate Change's Special Report on Emission Scenarios. Variables are available for a regular one-eighth degree grid covering the continental U.S., southern Canada, and northern Mexico, as well as for over 8,000 long-term weather stations in North America. The dataset consists of more than 6 terabytes of data and was generating using over 2 million CPUs of parallel computing time provided free of charge by the Texas Tech University High Performance Computing Center.

Continuous daily outputs for the period 1960 to 2099 were obtained from 11 CMIP3 global climate models (GCMs): NCAR-CCM3, CCCMA-CGCM3_T47, CCCMA-CGCM3_T63, CNRM, ECHAM5, ECHO, GFDL CM2.0, GFDL CM2.1, UKMO-HadCM3, UKMO-HadGEM1 and NCAR-PCM. Limited daily time slices were obtained for 5 additional models: CSIRO, GISS_AOM, MIROC (high-res), MIROC (med-res), and MRI_CGCM2. Not every model had daily outputs available for each emissions scenario; the total number of model/scenario combinations available is 48 for temperature and 45 for precipitation. Although the GCM outputs conformed to the standards of submission to CMIP3, during the research a number of errors were discovered in the original GCM files including missing data, erroneous data values, misnamed files, and (for one model) a consistent failure to simulate observed temperature and precipitation over the continental U.S. For that reason, **the project was expanded to include evaluation and quality control of all GCM inputs.**

Gridded daily meteorological data for the period 1960 to 2000 at one-eighth degree over the continental U.S. and one-half degree over Alaska were obtained from the database described in Maurer et al. (2002), and station data from the Global Historical Climatology Network (GHCN), enhanced by additional data records in sparse areas from the NWS-NCDC Cooperative Observer Network (COOP) and the U.K. Met Office Integrated Data Archive System (MIDAS). Although most station data had already been quality controlled, this research revealed that none of the quality control standards for these datasets were adequate for the purpose of downscaling. Specifically, the existence of even one anomalous extreme value could skew the statistical downscaling method, producing unreliable results for that location. Hence, **the project was expanded to include development of a rigorous quality control algorithm to pre-filter all station data.**

Initial downscaling using raw GCM output to the regular one-eighth degree grid revealed three flaws in downscaling methodology that had not been immediately apparent in previous limited-location studies:

1. Hundreds of years' worth of computing time was required to generate high-resolution climate projections for 75,000 grid cells, over 8,000 weather stations, 48 climate scenarios and 3 variables. This problem was addressed by professional optimization of

the downscaling code and parallelization with unlimited use of over 2 million CPUs from the 1024-core High Performance Computing Center at Texas Tech University.

2. Temperature fields from GCMs contained low-level noise that degraded the quality of the fit. This problem was addressed by pre-filtering GCM outputs using principal component analysis. The percentage of variance retained, 97 percent, was determined empirically by optimizing the fit of the statistical model. Filtered temperature fields were then used to downscale both gridded and station-based observations.
3. Downscaling grid cells independently for precipitation generated an unacceptable level of pixilation in individual rainy day fields for high precipitation events due to the sensitivity of the downscaling model at the tails of the distribution. This problem was addressed by: (a) modifying the precipitation downscaling routine to train the statistical model using a 3x3 grid surrounding each grid cell, thus ensuring spatial continuity at the extremes of the distribution, (b) spatially interpolating GCM output fields to the scale of the observations before using as predictors, and (c) as available, using GCM convective and large-scale precipitation outputs separately and together as predictors for individual grids.

After downscaling, the outputs were subjected to a comprehensive set of tests for 22 types of errors and 74 types of biases in the high-resolution climate projections. This quality control process clearly indicates that the downscaling method is capable of producing reliable information that closely replicates observations. At the same time, this analysis was able to quantify how the skill of the downscaling—particularly for extreme values—also depends on the performance of the original GCM and on the range of conditions incorporated by the historical observations used to train the downscaling model.

In addition to the bias and error analysis, 114 impact-relevant indicators, from seasonal temperature and precipitation to dry days and heat waves, were calculated for five time periods (1960-1979, 1990-2009, 2020-2039, 2050-2069, 2080-2099) as well as for global mean temperature intervals of +1°C, +2°C and +3°C relative to 1971-2000. Projected changes in temperature-related indicators, in heavy precipitation, and in dry days are all greater by end-of-century, under higher emissions, and for higher global mean temperature change as compared to earlier projections under lower emissions or lower global temperature change.

This research has resulted in the most extensively tested set of high-resolution climate projections available for scientific research in the U.S. The projections generated by this project will be used as input to the 2014 U.S. National Climate Assessment as well as a host of other regional projects around the country. **It has also contributed significantly towards understanding the process of and uncertainties involved in downscaling.** The knowledge of downscaling acquired during this project is already being used as input to CSC-funded research on statistical downscaling model inter-comparisons, as well as the National Climate Predictions and Projections (NCPP) Platform. Finally, **the computing framework established by this project is now being applied to downscale new IPCC AR5/CMIP5 GCM outputs** (in progress). As these results become available, they will be subjected to the same rigorous error and bias tests as this data, and outputs can be provided to the USGS GeoData Portal if desired.

PURPOSE AND OBJECTIVES

The purpose of this work was to create a comprehensive web-based dataset of high-resolution climate change projections that could be used to assess the impacts of climate change on ecosystems in the continental United States. The dataset was generated by applying advanced statistical downscaling methods to a comprehensive selection of global climate model simulations from the IPCC AR4/CMIP3 database. The work allows for consistent impact assessments at the scale of the most critical ecosystem processes through downscaling projections of daily temperature and precipitation across the continental U.S. It also enables scientists and managers to easily access, manipulate and download data relevant to modeling climate change impacts on ecosystems through a common web-based data portal. Most importantly, the dataset generated by this work allows impact assessments to be based on the same common data set, producing consistent results to be compared across regions and ecosystems, addressing NCCWSC Goals 1 and 4 by developing and providing the high-resolution climate change datasets necessary for ecosystem impact studies.

All the objectives of the original work were met and **substantially expanded** to address questions, concerns, and needs that arose in preliminary evaluation and application of the high-resolution projections. In addition to the work originally proposed and supported by the funded project, the following **additional tasks** were also accomplished:

1. A **new pre-processing approach** was developed to remove spurious noise in global climate model daily temperature fields and spatially interpolate all global model output fields prior to their use in statistical downscaling. *This approach can be used to improve inputs to any downscaling method in the future.* Its application significantly improved biases in downscaled simulations, particularly for cold temperatures and for spatial patterns of precipitation.
2. The ARRM downscaling model was modified to allow for **multiple precipitation predictors** (convective, large-scale, and total precipitation). This modification enabled the model to better resolve the drivers of precipitation in summer vs. winter and in warmer vs. colder climates, and reduced biases in seasonal and extreme precipitation.
3. After identification of pixilation problems with precipitation extremes, the ARRM downscaling model was modified to train gridded precipitation models on a 3x3 grid in order to stabilize the statistical fit at the tails of the distribution and all precipitation downscaling was re-run with the new code. This modification **significantly reduced high-frequency noise**, or pixilation, for individual storms, resulting in spatial statistics for downscaled data that closely resembled observed.
4. An analysis **code to assess the spatial dependence of daily precipitation fields** in terms of the visual appearance of rain storms, the distance between pixel values, and Moran's I, a measure of spatial coherence, was developed to evaluate model simulations. *This is a generalizable code that can be used to test and compare the coherence of any spatial data field.*
5. The computing framework for conducting the downscaling was professionally **optimized** by Dr. Matthew Pratola, an expert in R optimization, and parallelized to

enable application using the Texas Tech University High Performance Computing Center (TTU-HPCC) resources. *This optimized computing framework will be used for future downscaling of CMIP5 simulations.*

6. A rigorous **quality control and nearest-neighbor process** was developed to identify and remove erroneous observed temperature and precipitation values in the station-based observations. *This is a generalizable code that can be used to quality control any observational dataset.*
7. **Additional observational datasets** (MIDAS, NCDC) were obtained, processed, and incorporated into this dataset. MIDAS was not available in daily format; sub-daily data was extracted and fit to a diurnal function before being evaluated relative to WMO standards. MIDAS and all additional datasets were processed using the quality control framework in order to provide as many neighbors as possible for existing GHCN locations. *Long-term records from these datasets have also been downscaled and are now incorporated into the database of high-resolution projections, expanding the original geographic scope of the analysis to encompass the entire continent of North America.*
8. A **comprehensive error and bias analysis code** was built which evaluates 22 types of errors and 74 types of biases in the high-resolution climate projections relative to observations. *This is a generalizable code that can be used to quality control any downscaled dataset or to compare any observational datasets.*
9. A **secondary indicator calculation code** was written to calculate 114 seasonal and annual mean and extreme indicators relevant to impact analyses. These indicators include seasonal temperature and precipitation, multi-day events, heavy precipitation and drought extremes, high and low temperature extremes, growing season, degree-days, etc. *This is a generalizable code that can be used to calculate secondary indicators for any set of observational, global model, or downscaled data, gridded or for individual locations.*
10. **Future projections were analyzed** in terms of future time periods and projected changes per degree global mean temperature change. *Maps and data files for indicators have been provided as accompanying information to the database.*
11. A brief **users guide** to high-resolution climate projections was written after many user-reported “errors” in the data were found to be misinterpretations or misuse of the data (e.g., assuming that any given simulated day should match observed conditions on that day, or being unaware that missing values in the downscaled projections were the result of missing data in the GCM input fields). *This text provides general information relevant to the application and use of any climate model simulations in impact assessments.*

ORGANIZATION AND APPROACH

The research process of this project consists of three main stages: obtaining and processing the inputs and code; preparing and conducting the simulations; and analyzing the outputs. The order of and relationships between the three primary research tasks are summarized in Figure 1. As errors were discovered with inputs and simulations, these processes were repeated multiple times. Systematic errors and specific problems that were identified and corrected or addressed during this process are discussed further in **ANALYSIS AND FINDINGS**.

In terms of inputs, creating a database of high-resolution climate projections requires global climate model simulations and observations at the desired spatial and temporal scale. Here, we used GCM simulations from CMIP3 and both gridded and station-based observations from a range of sources (Figure 1, blue boxes). These data are described in

more detail in this section, below. This research also required a statistical downscaling model. We used a model capable of resolving daily extremes, the Asynchronous Regional Regression Model (ARRM; Figure 1, purple box). The historical and future emission scenarios, GCM inputs, observations, and downscaling model are all described in more detail in this section, below.

During the course of this project, it was determined that pre-processing the GCM inputs and quality-controlling the station-based inputs was essential to generating robust projections at the regional scale (Figure 1, red/orange boxes). As awareness of, and solutions to, these two steps are an important result of this project, the GCM pre-processing and observational quality control process are described in more detail in **ANALYSIS AND FINDINGS**.

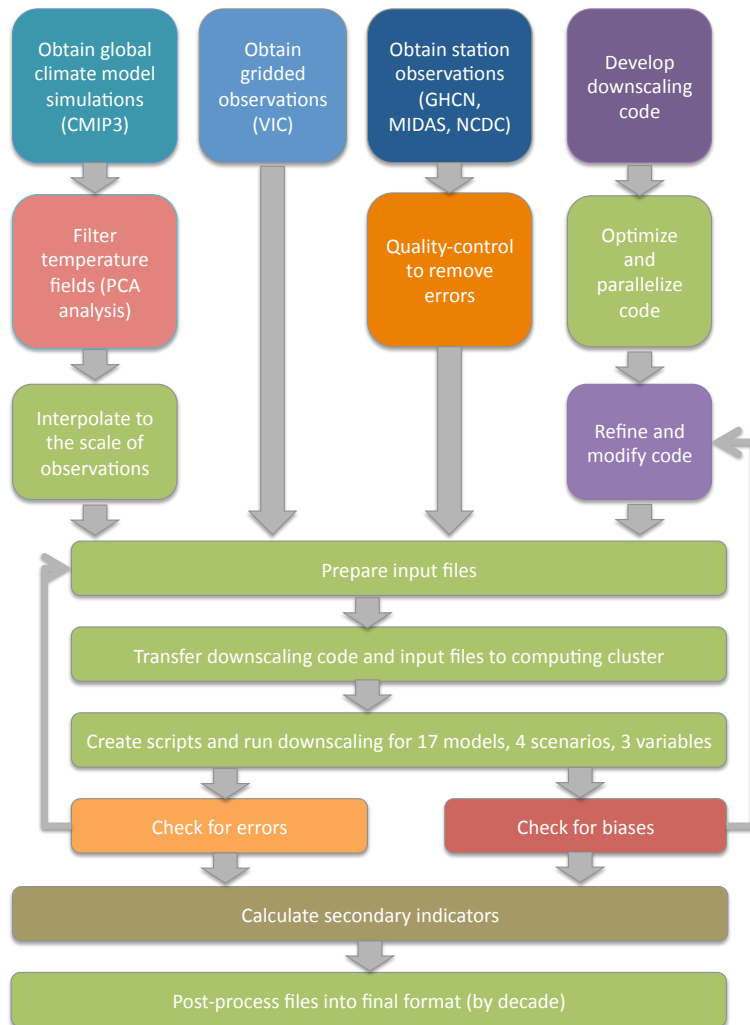


Figure 1. Schematic of the research tasks and activities conducted under this project.

Preparation of the code and input files and computational methods and post-processing of the resulting files (Figure 1, green boxes), although time intensive, can be described relatively briefly; these are also summarized in this section, below.

After downscaling, the projections were subjected to a rigorous quality control including checks for both errors and biases (Figure 1, red/orange boxes). As a result of these checks, a number of improvements were made to the downscaling code and input preparation. The error and biases tests and the method improvements are described in more detail in the section **ANALYSIS AND FINDINGS**.

Historical and Future Climate Scenarios

This project used SRES emission scenarios A1fi (higher), A2 (mid-high), A1B (mid-low) and B1 (lower). These scenarios were chosen because they cover a broad range of plausible futures in terms of human emissions of carbon dioxide and other radiatively-active species and resulting impacts on climate. The new RCP scenarios were not used yet because global climate model simulations using the RCP scenarios were not available until 2012. RCP-based projections will be generated and made available in the future if desired.

Climate model simulations begin with a long, multi-century “*control*” run where external forcing conditions including greenhouse gas concentrations, solar radiation, and volcanoes are fixed at constant levels corresponding to a specific year, generally in the 19th century. The choice of control year varies from one modeling group to the next, but is typically between 1850 and 1890. This long run is required for the ocean and atmospheric components of the model to equilibrate with each other and reach a stable climate. Output from control runs was not used in this project.

Once climate conditions are stabilized, the output from the control run can be used as input to a *transient historical simulation*. During a transient simulation, the external forcings (including greenhouse gas concentrations, solar radiation, and volcanic eruptions) change from year to year consistent with observed values for that year. The transient historical forcings used by the GCM simulations in this project are the Coupled Model Intercomparison Project’s “20th Century Climate in Coupled Models” or 20C3M total forcing scenarios (Meehl et al. 2007). These scenarios include forcing from anthropogenic or human emissions of greenhouse gases, aerosols, and reactive species; changes in solar output; particulate emissions from volcanic eruptions; changes in tropospheric and stratospheric ozone; and other influences required to provide a complete picture of the climate over the last century. As such, these simulations provide the closest approximation to actual climate forcing from the beginning of the transient experiment to the year 2000. Where multiple 20C3M simulations were available, the first was used here (“run 1”) unless complete daily outputs were not available for that simulation, in which case the next available was used.


In the same way as the control run can provide the starting conditions for multiple historical transient simulations, the historical transient simulation provides the starting conditions for multiple future simulations. To ensure the accuracy of the historical total forcing scenarios, it is customary in the climate modeling community for historical simulations to end at least 5 years before the actual year in which the simulation is being conducted. So although the CMIP3 GCM simulations were typically conducted after 2005,

CMIP3 historical total-forcing scenario ends and “future” scenarios begin in 2000. In the future scenarios, solar and volcanic forcing is fixed, and anthropogenic emissions correspond to a range of plausible pathways rather than observed values.

The scenarios used here are those described in the Intergovernmental Panel on Climate Change (IPCC) Special Report on Emissions Scenarios (SRES; Nakićenović et al., 2000). These scenarios describe internally consistent pathways of future societal development and corresponding emissions. The carbon emissions and global temperature change that result from the SRES scenarios are shown in Figure 2 (left side).

At the higher end of the range, the SRES higher-emissions or fossil fuel intensive scenario (A1FI, for *fossil-intensive*) represents a world with fossil fuel-intensive economic growth and a global population that peaks mid-century and then declines. New and more efficient technologies are introduced toward the end of the century. In this scenario, atmospheric CO₂ concentrations reach 940 parts per million by 2100, more than triple pre-industrial levels of 280 ppm. At the lower end, the SRES lower-emissions scenario (B1) also represents a world with high economic growth and a global population that peaks mid-century and then declines. However, this scenario includes a shift to less fossil fuel-intensive industries and the introduction of clean and resource-efficient technologies. Emissions of greenhouse gases peak around mid-century and then decline. Atmospheric carbon dioxide levels reach 550 parts per million by 2100, about double pre-industrial levels. Associated temperature changes by end-of-century range from 4 to 9°F based on the best estimate of climate sensitivity.

For this project, climate projections were based on the A1FI higher (dark red), A2 mid-high (orange), A1B mid-low (red) and B1 (blue) lower scenarios. Due to the decision of IPCC Working Group 1 to focus on the A2, A1B and B1 scenarios, only four GCMs had A1FI scenarios available. For other models, daily outputs were not available for all scenarios. Table 2, in the next section on **Global Climate Model Simulations**, summarizes the combinations of GCM simulations and emission scenarios used in this work.



In 2009, the IPCC released a new set of scenarios, called *Representative Concentration Pathways* or RCPs (Moss et al., 2010). In contrast to the SRES scenarios, the RCPs are expressed in terms of carbon dioxide concentrations in the atmosphere, rather than direct emissions. However, RCP scenarios can be converted “backwards”, into the range of emissions consistent with a given concentration trajectory, using a carbon cycle model (Figure 2, center). Four RCP scenarios were developed to span a plausible range of future carbon dioxide concentrations, from lower to higher. At the higher end of the range, atmospheric carbon dioxide levels under the RCP 8.5 scenario reaches more than 900 parts per million by 2100. At the lower end, under RCP 2.6 policy actions to reduce carbon dioxide emissions below zero before the end of the century (i.e. to the point where humans are responsible for a net uptake of carbon dioxide from the atmosphere) keeps atmospheric carbon dioxide levels below 450 parts per million by 2100. Associated temperature changes by end-of-century range from 2 to 8°F based on the best estimate of climate sensitivity. The RCP scenarios were not used in this work because, although the scenarios themselves were released in 2009, global climate model simulations using the RCP scenarios as input were not available until 2012. The SRES scenarios were used because those were the ones on which available CMIP3 GCM outputs were based.

The implications of using the SRES scenarios instead of RCP scenarios can be quantified by comparing the scenarios’ emissions (Figure 2) and carbon dioxide concentrations, or atmospheric levels (Table 1). There are clear RCP-SRES “pairs”: i.e., scenarios that closely resemble each other in terms of climate forcing and which would therefore be expected to result in a similar amount of climate change. These are the higher SRES A1fi and RCP 8.5; mid-low SRES A1B and RCP 6.0; and lower SRES B1 and RCP 4.5. Since the SRES scenarios begin in 2000, they can also be evaluated relative to historical emissions over the last decade. This comparison shows that actual emissions are currently near the top of the range of both SRES and RCP scenarios (Figure 2b).

Table 1. Comparing SRES emission and RCP concentration scenarios, ranked from high to low in order of carbon dioxide levels in the atmosphere (in units of parts per million by volume) by 2100.

Scenario name	CO2 in 2100	What to call these scenarios?
SRES A1fi	958	Higher emissions or fossil-fuel intensive
RCP 8.5	936	
SRES A2	836	Mid-high emissions or continued emissions growth
SRES A1b	703	Higher mid-range; emissions peak, then begin to decline
RCP 6.0	670	
SRES B2	611	Lower mid-range; emissions peak, then decline
SRES A1T	575	
SRES B1	540	Lower emissions; emissions peak mid-century, then decline to near 1990 levels by 2100
RCP 4.5	538	
RCP 2.6	421	Stringent mitigation scenario Negative emissions before end of century

Global Climate Model Simulations

This project used CMIP3 global climate model simulations from 16 different models. These include all but one of the CMIP3 models that had archived the daily maximum and minimum temperature and precipitation outputs required for downscaling. The new CMIP5 global climate model outputs were not used because they were not available until 2012 and are still being updated in 2013.

Global climate model simulations, while in a state of constant flux within a given research group or laboratory, are archived at regular intervals by the World Climate Research Programme's Working Group on Coupled Modelling (WGCM). In preparation for the IPCC's Fourth Assessment Report (AR4), the WGCM requested that the US DOE-funded Program for Climate Model Intercomparison and Diagnosis (PCMDI) collect model output from climate modeling centers around the world. This first collection, assembled between 2005 and 2006 and archived by PCMDI, represents models that contributed to phase 3 of the Coupled Model Intercomparison Project (CMIP3; Meehl et al., 2007). These are the results presented in the 2007 IPCC Fourth Assessment Report (AR4).

The GCM simulations used in this project consist of all CMIP3 simulations archived by PCMDI with daily maximum and minimum temperature and precipitation outputs. Additional simulations were obtained from the archives of the Canadian Centre for Climate Modeling and Analysis, the Geophysical Fluid Dynamics Laboratory, the National Center for Atmospheric Research, and the U.K. Meteorological Office. A total of 17 GCMs met this data-based criteria. The full list of GCMs used, their origin, the scenarios available for each, and the time periods covered by their output are given in Table 2. Output from 12 GCMs was available for the full time period (1960 or 1961 to 2099) while output from 5 more GCMs was available for three time slices (1961-2000, 2046-2065, 2081-2100).

After the original GCM files were obtained from their host archive, they were subjected to a basic quality control to ensure the files contained the days and the data they stated that they did, that the data was within reasonable bounds for the variable listed, and that any missing data were identified. This analysis showed that many original GCM files had errors or peculiarities that were catalogued by this project before conducting the downscaling. The results of the quality control process and the errors it detected are summarized in **ANALYSIS AND FINDINGS**.

No attempt was made to select a sub-set of GCMs that performed better than others, as previous literature has showed that it is difficult, if not impossible, to identify such a sub-set for the continental U.S. (e.g. Knutti, 2010; Randall et al. 2007) However, the bias and error analysis conducted here offers insights into model performance (see **PROJECT RESULTS**), including demonstrating how some models are better able to simulate regional aspects of temperature and precipitation over the continental U.S. than others and identifying one model (BCCR-BCM2) with consistently poor performance that was removed from the archive. BCCR-BCM2 is a relatively new model and does not have a more recent version included in the CMIP5 phase, as do the remainder of the models used here.

From 2011 through the end of 2012, PCMDI began to collect and archive new GCM simulations that contributed to the fifth phase of CMIP and which will be used in the upcoming IPCC Fifth Assessment Report (AR5; Taylor et al. 2012). The CMIP3 and CMIP5

archives are similar in that most of the same international modeling groups contributed to both. There are even a few of the same models (e.g. HadCM3) in both CMIP3 and CMIP5. Both provide daily, monthly, and yearly output from climate model simulations driven by a wide range of future scenarios. However, the archives are also different from each other in three key ways. First, many of the CMIP5 models are new versions or updates of previous CMIP3 models and some of the CMIP5 models are entirely new. Some of the CMIP5 models are “Earth System Models” that include both traditional components of the CMIP3 Atmosphere-Ocean General Circulation Models as well as new components such as atmospheric chemistry or dynamic vegetation. Second, the CMIP5 simulations use the RCP scenarios as input for future simulations while the CMIP3 simulations use the SRES scenarios as input (Figure 2). Third, the CMIP5 archive contains many more output fields than the CMIP3 archive did. **Future downscaling efforts will apply the same methods and computational framework to CMIP5 simulations, so these can be used in conjunction with the CMIP3 simulations developed here, as a “super-ensemble”.**

Table 2. Global climate modeling groups and their models, output from which they have contributed to CMIP3 (used in IPCC AR4, 2007) and/or CMIP5 (used in IPCC AR5, 2013). Those marked with (*) archived only time slices: 1961-2000, 2046-2065 and 2081-2100. Those marked with (^) begin in 1961. All other models archived full daily time series from 1960 to 2099. The BCCR-BCM2 model was eliminated from the archive due to poor performance.

Origin	CMIP3 model(s)	Scenarios
Bjerknes Centre for Climate Research, Norway	BCCR-BCM2.0	A2, A1B, B1
National Center for Atmospheric Research, USA	CCSM3	A1FI, A2, A1B, B1
	PCM	A1FI, A2, A1B, B1
Canadian Centre for Climate Modelling and Analysis, Canada	CGCM3.1 – T47 [^]	A2, A1B, B1
	CGCM3.1 – T63 [^]	A2, A1B, B1
Centre National de Recherches Meteorologiques, France	CNRM-CM3	A2, A1B, B1
Commonwealth Scientific and Industrial Research Organisation, Australia	CSIRO-Mk3.0* [^]	A2, A1B, B1
Max Planck Institute for Meteorology, Germany	ECHAM5/MPI	A2, A1B, B1
National Institute of Meteorological Research/Korea Meteorological Administration	ECHO-G (with MPI)	A2, A1B, B1
NOAA Geophysical Fluid Dynamics Laboratory, USA	GFDL CM2.0	A1FI, A2, A1B, B1
	GFDL CM2.1	A2, B1
NASA Goddard Institute for Space Studies, USA	GISS-AOM* [^]	A1B, B1
UK Meteorological Office Hadley Centre	HadCM3	A1FI, A2, A1B, B1
	HadGEM1	A2, A1B
Agency for Marine-Earth Science and Technology, Atmosphere and Ocean Research Institute, and National Institute for Environmental Studies, Japan	MIROC3.2 (hires)* [^]	A1B, B1
	MIROC3.2 (medres)* [^]	A2, A1B, B1
Meteorological Research Institute, Japan	MRI-CGCM2.3.2* [^]	A2, A1B, B1

Downscaling Method

This project used the statistical Asynchronous Regional Regression Model (ARRM). It was selected because it is able to resolve the tails of the distribution of daily temperature and precipitation to a greater extent than the more commonly used Delta and BCSD methods, but is less time-intensive and therefore able to generate more outputs as compared to a high-resolution regional climate model.

Dynamical and statistical downscaling represent two complimentary ways to incorporate higher-resolution information into GCM simulations in order to obtain local to regional-scale climate projections. Dynamical downscaling, often referred to as regional climate modeling, uses a limited-area, high-resolution model to simulate physical climate processes at the regional scale, with grid cells typically ranging from 10 to 50km per side. Statistical downscaling models capture historical relationships between large-scale weather features and local climate, and use these to translate future projections down to the scale of any observations—here, both individual weather stations as well as a regular grid.

Statistical models are generally flexible and less computationally demanding than regional climate models, able to use a broad range of GCM inputs to simulate future changes in temperature and precipitation for a continuous period from 1960 to 2100. Hence, statistical downscaling models are best suited for analyses that require a range of future projections that reflect the uncertainty in emission scenarios and climate sensitivity, at the scale of observations that may already be used for planning purposes. If the study is more of a sensitivity analysis, where using one or two future simulations is not a limitation, or if it requires multiple surface and upper-air climate variables as input, and has a generous budget, then regional climate modeling may be more appropriate.

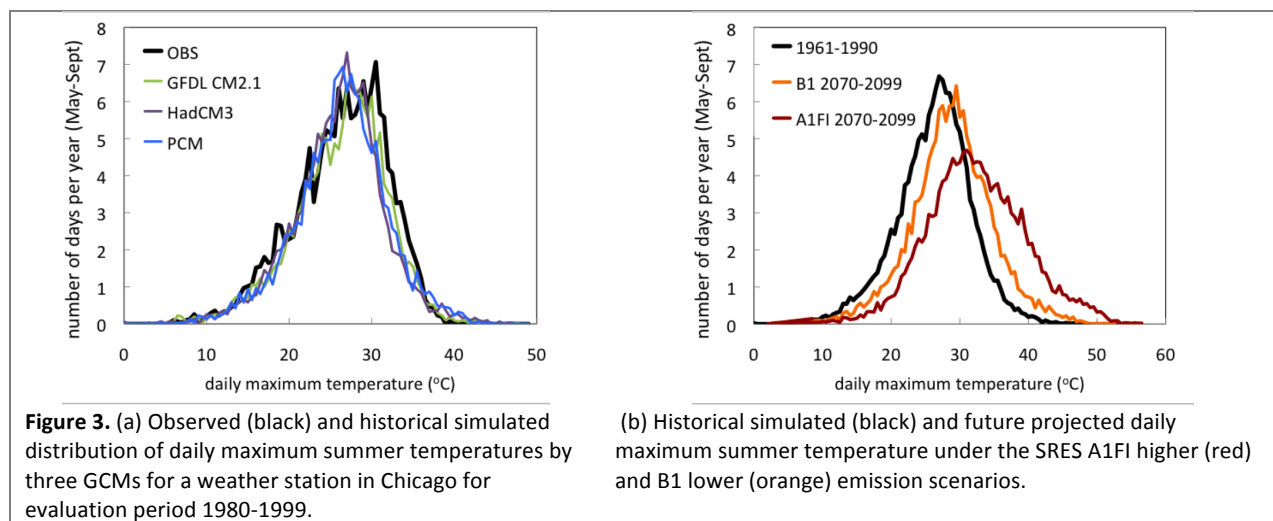
Each commonly used downscaling method has its own benefits, and each can be sufficient for certain applications. For example, the simple delta or “change factor” approach does a good job with downscaling annual or seasonal mean temperature (as demonstrated in Figures 4 and 5). Regional climate models require large amounts of computing power, but provide consistent high-resolution projections for a broad range of surface and upper-air variables. None of these existing methods, however, allow for using multiple climate models and scenarios as input while downscaling to any spatial scale (including both station-based and gridded), and adequately resolving projected changes in daily climate extremes, at the same time.

For that reason, in this project we used a relatively new statistical downscaling model, the Asynchronous Regional Regression Model, or ARRM (**Appendix F**: Stoner et al., 2012). ARRM uses asynchronous quantile regression, originally developed by Koenker and Bassett (1978) to estimate conditional quantiles of the response variable in econometrics. Dettinger et al. (2004) was the first to apply this statistical technique to climate projections to examine simulated hydrologic responses to climate variations and change, as well as to heat-related impacts on health (Hayhoe et al., 2004).

ARRM expands on these original applications with modifications specifically aimed at improving the ability of the model to simulate the shape of the distribution including the tails, including use of a piecewise rather than linear regression to accurately capture the

often non-linear relationship between modeled and observed quantiles, and bias correction at the tails of the distribution.

ARRM is a flexible and computationally efficient statistical model that can downscale station-based or gridded daily values of any variable that can be transformed into an approximately symmetric distribution and for which a large-scale predictor exists. A quantile regression model is derived for each individual grid cell or weather station that transforms historical model simulations into a probability distribution that closely resembles historical observations (Figure 3a). This model can then be used to transform future model simulations into distributions similar to those observed (Figure 3b). More information on the ARRM method is provided in **Appendix F**, “An asynchronous regional regression model for statistical downscaling of daily climate variables” by Stoner et al. (2012).

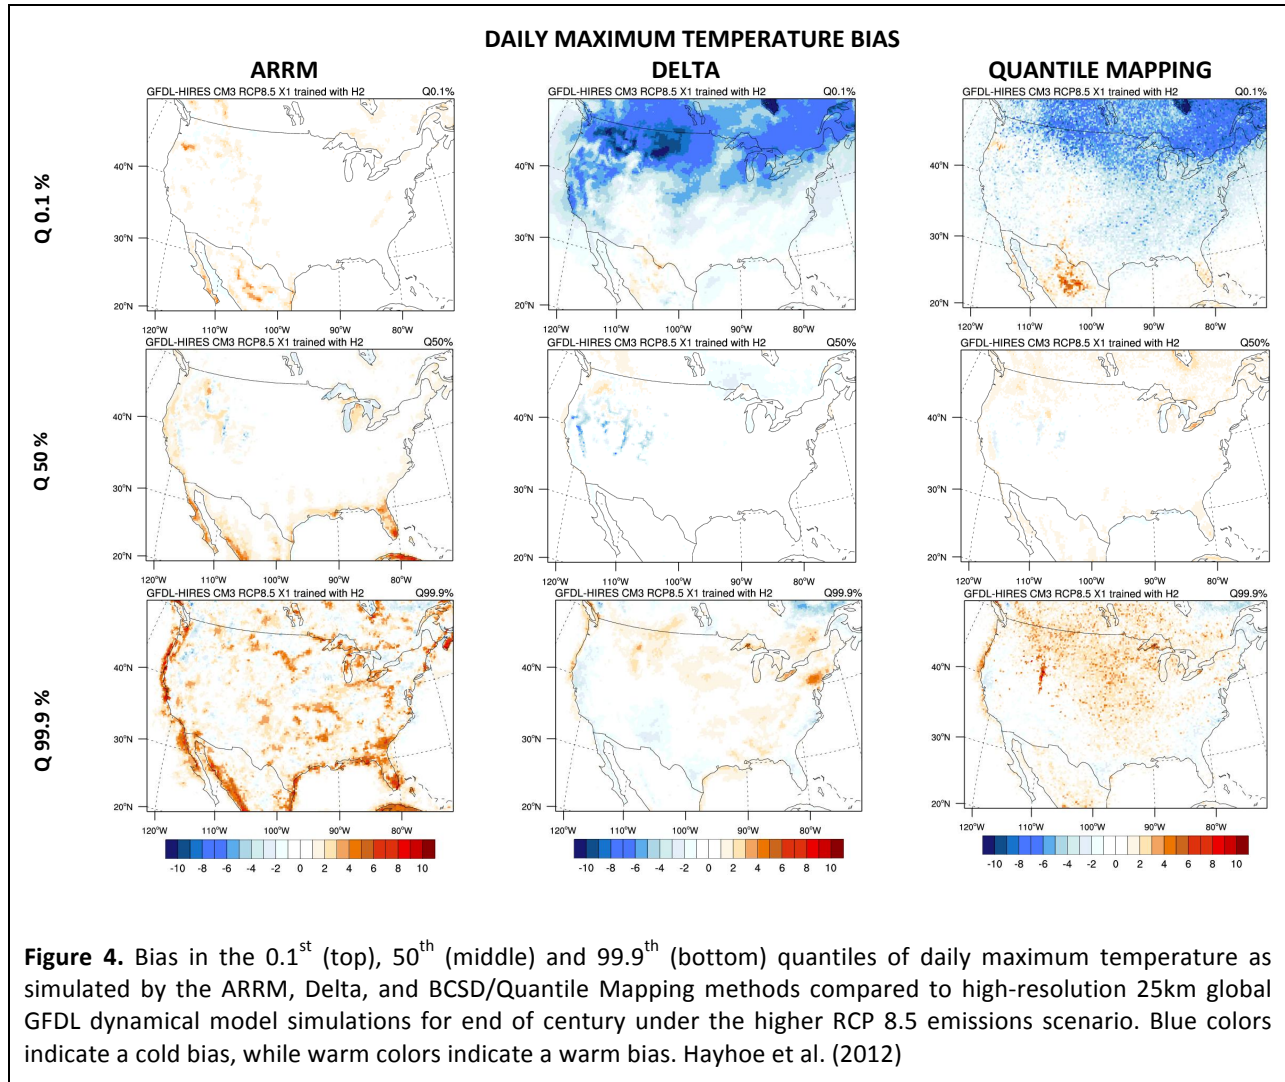


Both statistical and dynamical downscaling models are based on a number of assumptions, some shared, some unique to each method. Two important shared assumptions are the following: first, that the inputs received from GCMs are reasonable, i.e. that they adequately capture the large-scale circulation of the atmosphere and ocean at the skillful scale of the global model; and second, that the information from the GCM fully incorporates the climate change signal over that region.

All statistical models are based on a crucial assumption often referred to as **stationarity**. Stationarity assumes that the relationship between large-scale weather systems and local climate will remain constant over time. This assumption may be valid for lesser amounts of change, but could lead to biases under larger amounts of climate change (Vrac et al., 2007).

In a separate USGS-funded TTU-GFDL project, we are currently evaluating the stationarity of three downscaling methods, including the ARRM method (used here), the delta approach, which adds a “delta” derived from GCM output to observed mean annual, seasonal, or monthly values in order to get future values (e.g., Hay et al., 2000; as used in USGCRP, 2000); and the Bias Correction-Statistical Downscaling (BCSD) model, which uses a

quantile mapping approach to downscale monthly AOGCM-based temperature and precipitation to a regular grid (Wood et al. 2004; as used in Hayhoe et al. 2004, 2008 and USGCRP, 2009; available from ClimateWizard.org and the DOE Green Data Portal). In this ongoing project, high-resolution 25km global model simulations for end-of-century under the higher RCP 8.5 scenario have been coarsened and used as input to these three statistical downscaling methods. The resulting projections are then compared directly to the high-resolution output to determine the extent to which the assumption of stationarity holds true. Where biases are small, stationarity is a reasonable assumption. Where biases are large, the assumption of stationarity fails.



To examine the stationarity in daily maximum temperatures, Figure 4 compares biases from the 0.1th to the 99.9th quantile (i.e. from the coldest day in 1000 to the hottest day in 1000). These biases represent the difference between daily maximum temperature values simulated by the 25km global climate model vs. those simulated by each statistical model, using coarsened global climate model fields as input.

Preliminary analyses show that the assumption of stationarity holds true over much of the continental U.S. for the middle of the distribution (the 50th quantile, or median maximum temperature). The delta and BCSD methods exhibit a strong cold bias of up to 10°C at higher latitudes for cold temperatures. The BCSD method also begins to exhibit a systematic warm bias across the central U.S. at high temperatures. The only location where ARRM performance is systematically non-stationary is at high temperatures (at and above the 99.9th quantile) along the coast, with warm biases up to 6°C. This may be due to the statistical model's inability to capture dynamical changes in the strength of the land-sea breeze as the temperature differences between land and ocean are exacerbated under climate change; the origins of this feature are currently under investigation.

For minimum temperatures, Figure 5 compares biases from the 0.1th to the 99.9th quantile (i.e. from the coldest night in 1000 to the hottest night in 1000). Again, these biases represent the difference between nighttime minimum temperatures simulated by the 25km global climate model vs. those simulated by each statistical model from coarsened global climate model output. Here, the cold bias in monthly Delta and BCSD/Quantile Mapping methods is more widespread but less spatially consistent. For high minimum temperatures (warm nights), the ARRM method has a larger bias across much of the U.S. between 2 to 6°C, suggesting that it may over-estimate the increase in warm night temperatures in the future.

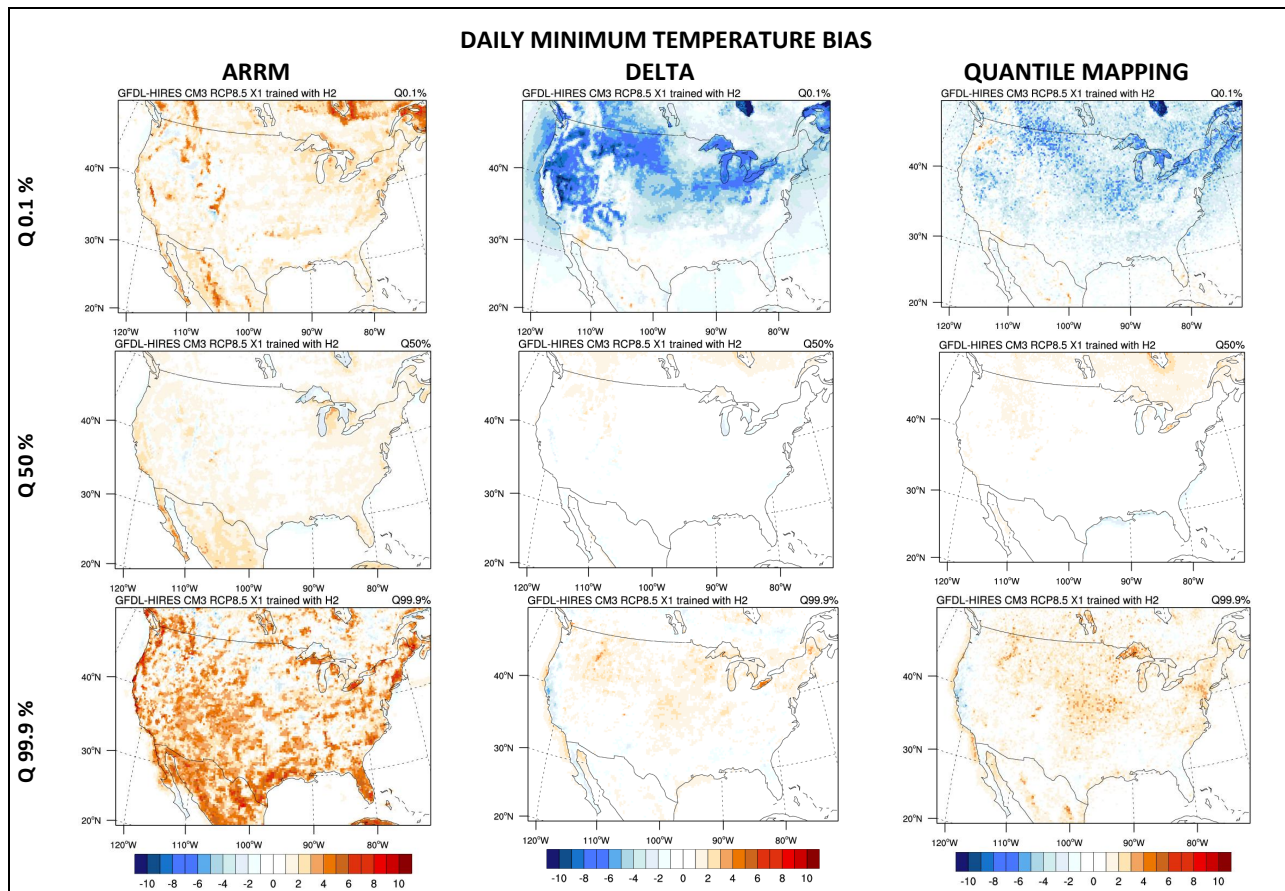


Figure 5. Same as Figure 3, but for minimum daily temperature. Hayhoe et al. (2012)

For precipitation, Figure 6 compares biases from the 10th to the 99.9th quantile, including both wet and dry days. Green colors are used to indicate where the statistical models over-estimate precipitation relative to the global model, while brown colors show where they under-estimate future precipitation.

Here, the BCSD approach, originally designed for forecasting streamflow that typically depends on accumulated precipitation over timescales of weeks (i.e. not high quantile events), shows a remarkable near-zero bias up to the 90th quantile. After that point, however, it rapidly develops a systematic positive bias in precipitation covering the entire continental U.S. but greatest in the Southeast region, increasing to 100% for the 99.9th quantile. (Recall that this is the bias in the entire distribution including dry days; hence, a bias in the 99th quantile is the bias in the 1-out-of-100 wettest day including both wet and dry days.) The Delta method exhibits a systematic positive (wet) bias for low precipitation that shifts to a systematic negative (dry) bias for high precipitation. The ARRM method is characterized by a spatially variable bias at all quantiles that *is generally not systematic*, and varies from approximately -30 to +30% depending on location. The magnitude of the bias increases for higher quantiles and appears to intensify somewhat around the Gulf Coast region by the 99.9th quantile, suggesting once more that the assumption of stationarity may not be able to capture the effects of land/ocean changes on coastal climate.

Although the downscaling model is purposely designed to be applicable to any variable with a relatively symmetric distribution, predictors must be pre-selected for each variable and there are some differences in the initial processing of each predictor that can improve the performance of the model in downscaling. The ARRM method has been specifically designed to allow for user-selected predictors. For the sake of consistency, however, in this project predictors were chosen to be the same variables as the predictands: 2m maximum and minimum temperature and 24h cumulative precipitation. These are the most frequently-archived daily output from both CMIP3 and CMIP5 AOGCMs; furthermore, comparison with upper-air predictors for the stations in this study showed no consistent continent-wide improvement that would affect the performance of the downscaling model.

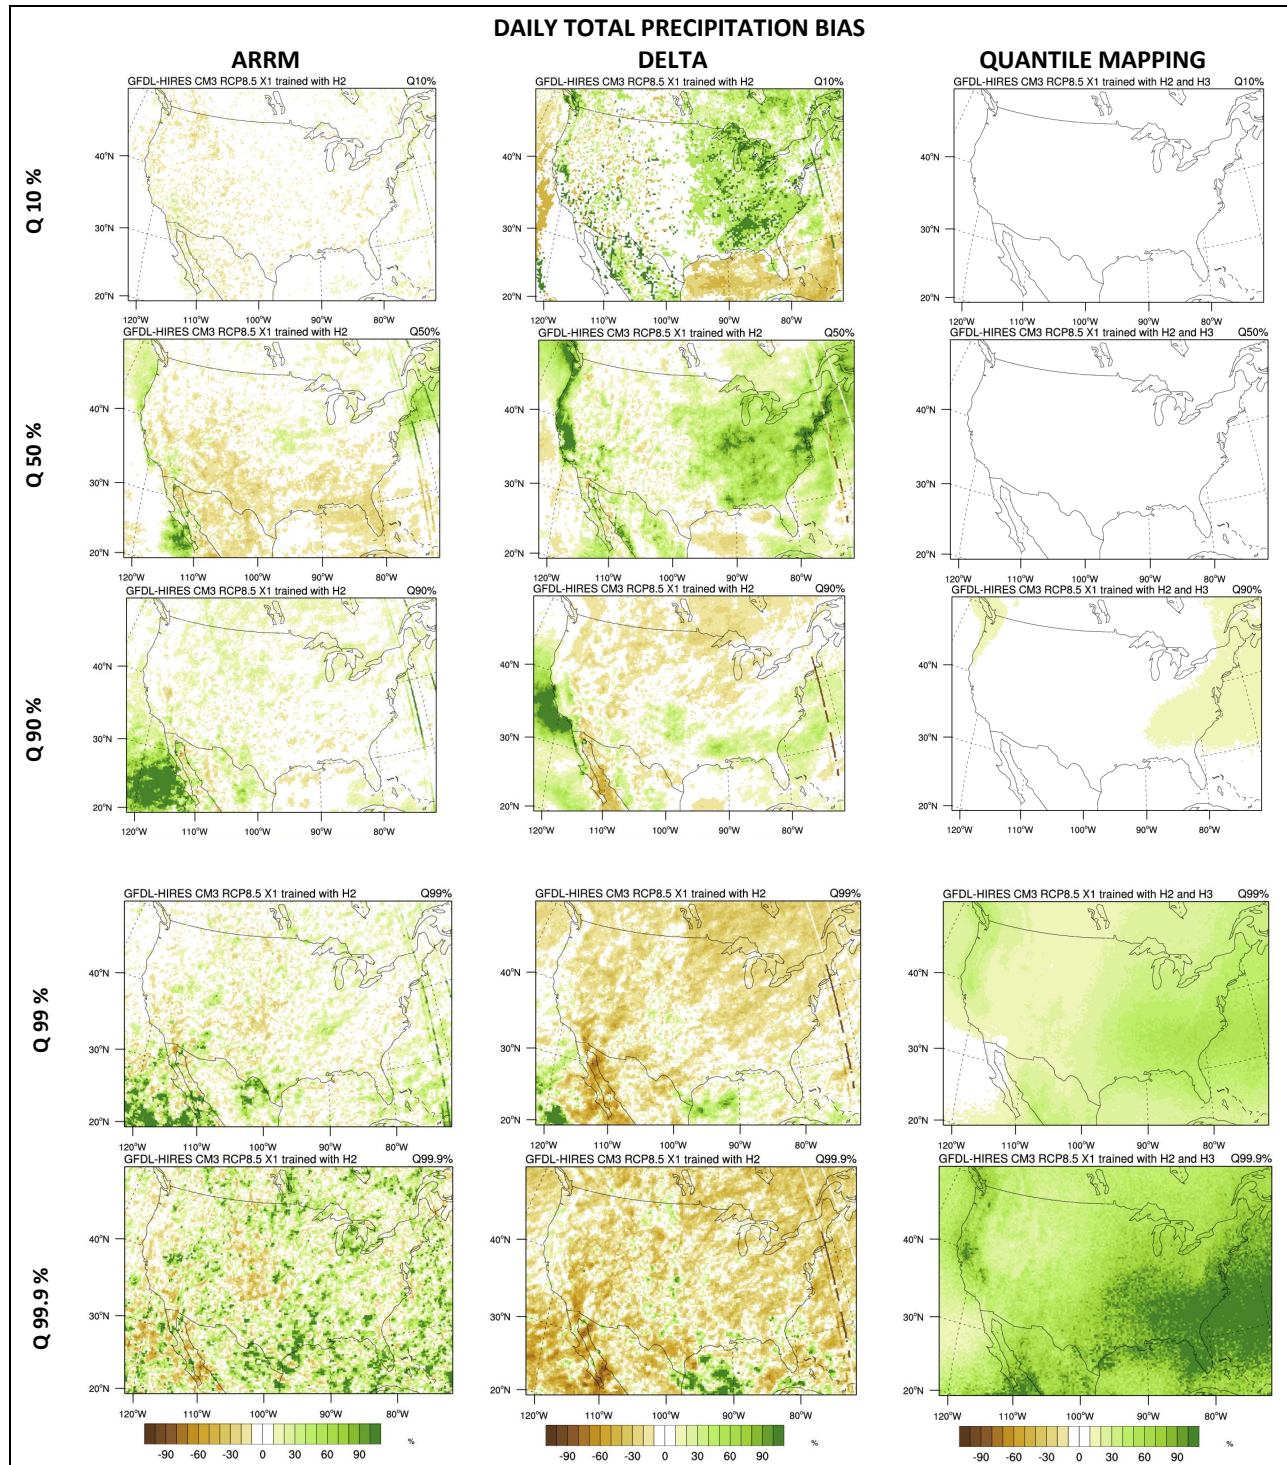
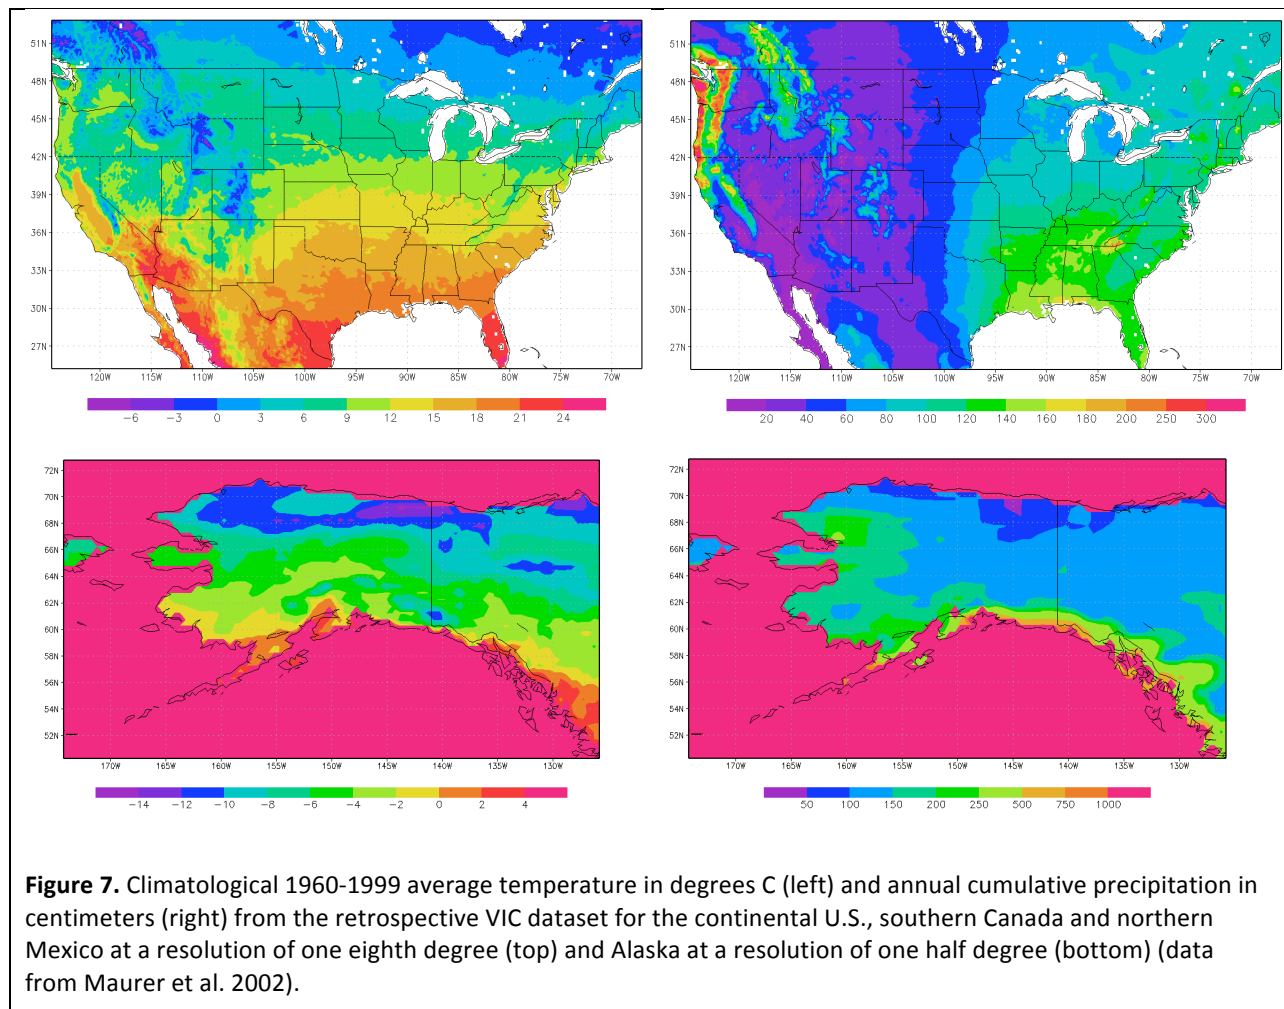


Figure 6. Bias in the 10th (first row), 50th (second row), 90th (third row), 99th (fourth row) and 99.9th (last row) quantiles of daily precipitation (including wet and dry days) as simulated by the ARRM, Delta, and BCSD/Quantile Mapping methods compared to high-resolution 25km global GFDL dynamical model simulations for end of century under the higher RCP 8.5 emissions scenario. Green colors indicate a wet bias, while brown colors indicate a dry bias. Hayhoe et al. (2012)

Gridded Observations

This project used the gridded meteorological data prepared by Maurer et al. (2002) at one-eighth degree over the continental U.S. and one-half degree over Alaska. At the time of the project, this was the longest, best quality, and highest-resolution dataset of gridded daily temperature and precipitation available. This original dataset has now been updated to a one-sixteenth grid; future work will use the updated dataset.

Daily maximum and minimum temperature and 24h cumulative precipitation from 1960 to 1999 for a regular grid covering the continental U.S., southern Canada and northern Mexico was used to train the statistical downscaling model, ARRM, for gridded downscaling. Over the continental U.S., the resolution of the grid is one-eighth degree, which can range from approximately 10 to 15km, depending on latitude. Over Alaska, the resolution of the grid is one-half degree. The resolution of the datasets are illustrated by the climatological 1960-1999 mean temperature and precipitation for the continental U.S. region and for Alaska, shown in Figure 7.



This data was obtained from the long-term hydrological dataset originally developed by and described Maurer et al. (2002) for long-term stream forecasting. As the authors note (see http://www.engr.scu.edu/~emaurer/gridded_obs/index_gridded_obs.html), this dataset has been checked to ensure the values are able to reproduce observed monthly

total precipitation and average temperature for randomly-selected grid cells only. For that reason, this gridded dataset was subjected to the same error tests and quality control process as the downscaled projections; the results of these tests, presented in Appendix C, suggest that this dataset generally free of obvious errors and thus provides a reliable basis for training the statistical model.

This original dataset has now (2013) been revised to cover the period 1915 to 2011 and the resolution increased to one-sixteenth degrees over the continental U.S. (Livneh et al. 2013; available at <ftp://ftp.hydro.washington.edu/pub/blivneh/CONUS/nc.1915.2011.bz2/>). Future downscaling using CMIP5 models will be based on the updated dataset, using the period 1950-2011 to train the statistical model and generate output fields at 1/16th degree resolution.

Station observations

This project used long-term station data from the Global Historical Climatology Network, supplemented with additional station data from MIDAS and NCDC. All station data was quality-controlled to remove questionable data points before being used to train the statistical downscaling model.

In addition to the gridded dataset, point-location weather station data was also obtained and used to train the downscaling model in order to generate a separate, independent point-based dataset of climate projections corresponding to specific weather stations. **These projections are particularly useful for locations with high spatial variability where a one-eighth degree grid may provide too much smoothing of local climate conditions, and/or for applications (such as streamflow or crop models) where impact models have already been calibrated to weather observations from specific weather stations.** For these applications, the station-based climate projections can be extracted and applied without requiring significant modification of the original experimental methodology.

Records were obtained from the Global Historical Climatology Network (GHCN; <http://www.ncdc.noaa.gov/oa/climate/ghcn-daily/>) and supplemented with additional records from the U.K. Met Office Integrated Data Archive System Land and Marine Surface Stations Data (MIDAS, <http://badc.nerc.ac.uk/data/ukmo-midas/>) and the National Climatic Data Center cooperative observer program (NCDC-COOP, <http://www.ncdc.noaa.gov/land-based-station-data/cooperative-observer-network-coop>). Additional station data used in previous projects (e.g. from the Arizona Mesonet and the Florida Climate Center) were also incorporated into the quality control and downscaling process to maximize the number of available data records, *if records for their latitude and longitude locations were not already present in GHCN or MIDAS*.

MIDAS was not available in daily format; thus, three-hour instantaneous temperature data was extracted and fit to a diurnal function and cumulative precipitation was extracted and summed before being evaluated relative to WMO standards (e.g. maximum daily temperature must occur between 6am and 6pm; minimum nighttime temperature between 6pm and 6am; precipitation is summed from 0h to 21h local time).

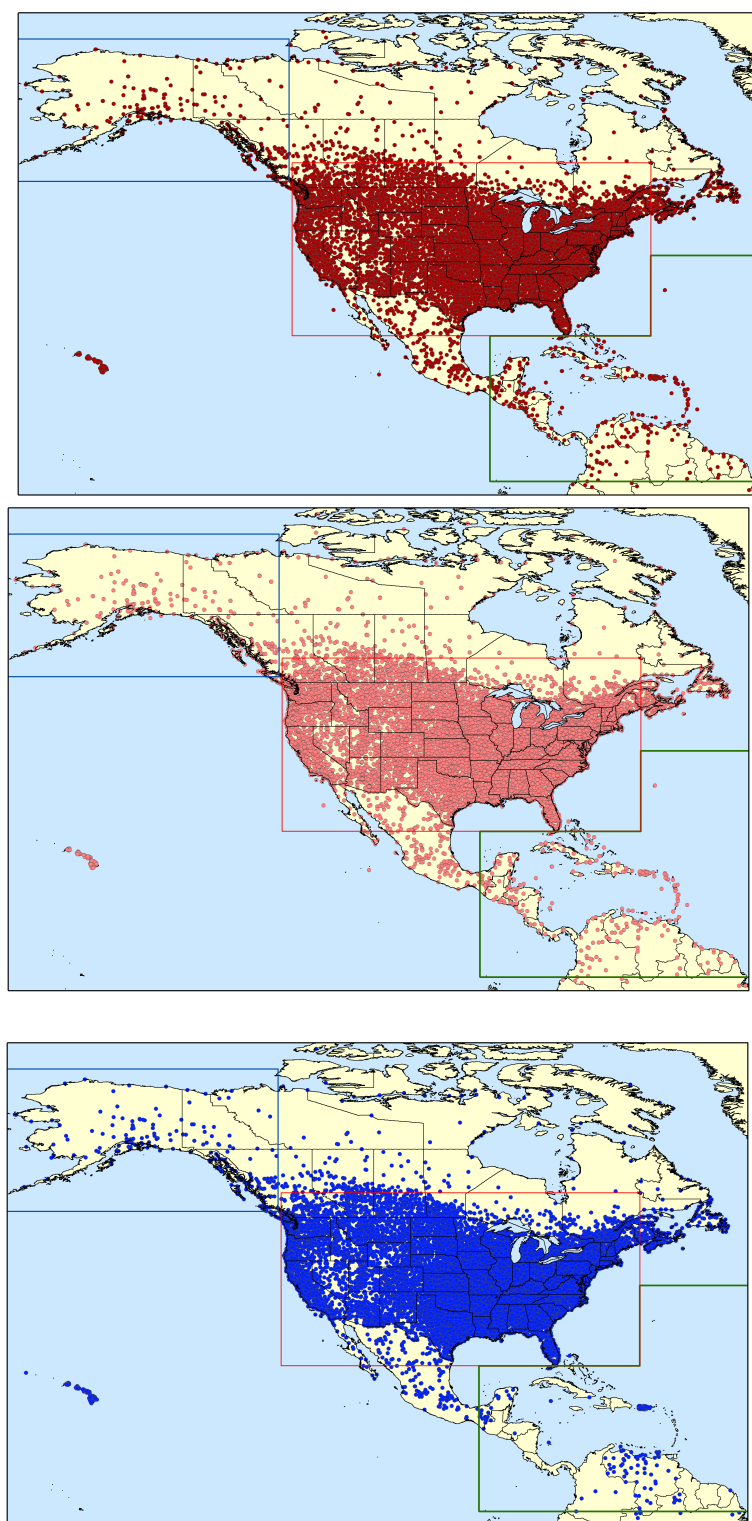


Figure 8. Maximum temperature (top), minimum temperature (middle) and precipitation (bottom) weather stations for Alaska (blue box), the continental U.S. (red box) and the Caribbean region (green box).

For downscaling, the weather stations were grouped into three geographic areas as shown in Figure 8: Alaska (blue box), continental U.S., southern Canada and northern Mexico (red box) and U.S. Caribbean territories (green box). Stations outside these regions (specifically, locations in north central and eastern Canada and southern Mexico) were not downscaled.

To train the downscaling model, the observed record must be of adequate length and quality. To appropriately sample from the range of natural climate variability at most of the station locations, and to produce robust results without over-fitting, stations were required to have a minimum of 20 consecutive years of daily observations overlapping GCM outputs with less than 50% missing data after quality control. When these limits were applied, the number of usable stations was reduced from 10,986 to 8,212 for maximum temperature, from 10,920 to 8,176 for minimum temperature, and from 13,233 to 10,272 for precipitation. The latitude, longitude, and station names of the weather stations for which downscaled projections were generated are provided in **Appendix A**.

All datasets were incorporated into the quality control framework described in **ANALYSIS AND FINDINGS**.

Long-term records from these datasets were then downscaled and incorporated into the database of high-resolution projections described in **PROJECT RESULTS**, expanding the original geographic scope and resolution of the grid-based downscaling.

Computational Methods

Data file formats and computational approaches used represent the standard for climate model data and analysis. Datasets are archived as ascii text files (observed station data) and self-describing netCDF files (gridded observations and both station-based and gridded climate projections). All codes were written using the statistical programming language R (<http://cran.us.r-project.org/>). Figures were prepared using the Grid Analysis and Display System (<http://www.iges.org/grads/>) and the NCAR Command Language (<http://www.ncl.ucar.edu/>) and stored in postscript or Adobe PDF format.

PROJECT RESULTS

This project produced four quantitative datasets, as follows:

ONE. A quality-controlled, standardized database of long-term observed daily maximum and minimum temperature and 24h cumulative precipitation for over 8,000 individual long-term weather station locations across North America.

This dataset was generated from original station data, covering the period of record for each station. Potentially erroneous data points were removed during the quality control process described in **ANALYSIS AND FINDINGS**. The resulting data was stored in individual comma-separated ASCII files containing four columns: year, month, day, and value. Individual files were created for each location and variable, with the files named by the station ID. A master list was also created for each variable, containing the latitude, longitude and ID for each station. This list was used to eliminate duplicate records across multiple databases.

Stations with sufficient data availability for downscaling that passed the quality control process are listed in **Appendix A**.

This dataset is not being provided electronically at this time. Although it is a derivative of existing datasets, it contains enough non-original data that additional permissions (from GHCN, NCDC, MIDAS, and other sources) will be required before it can be archived and made available to users. Our permission is limited to individual use.

TWO. A database of daily maximum and minimum temperature and 24h cumulative precipitation projections for a regular one-eighth degree grid covering the land area of the continental U.S., southern Canada, and northern Mexico and a one-half degree grid covering Alaska.

This dataset was generated by downscaling global climate model outputs to land-based observations at the resolution of one-eighth degree for the continental U.S., southern Canada and northern Mexico (222 latitude bands and 462 longitude bands for a total of 102,562 individual gridcells per file, approximately 25% of which are ocean) and one-half

degree for Alaska (46 latitude bands by 98 longitude bands for a total of 4,508 grid cells per file, approximately 30% of which are ocean).

Daily historical simulations and future projections were generated for 16 global climate models and 4 future SRES scenarios for the period 1960 to 2099 for a total of 51,100 days per file. The resulting data was archived in self-describing netCDF files, one for each variable and climate model/scenario combination. NetCDF format was selected for three reasons. First, this format allows individual locations and regions to be easily extracted using NCO or CDO command-line utilities. Second, this format provides a relatively efficient way of storing time-dependent spatial data. Lastly, this format is easily accessible by many commonly-used programs, including ArcGIS and R.

As the files are approximately 42 GB each, these are provided in electronic format only.

The format of the files is as follows (sample file header for the continental U.S.):

```
netcdf hadcm3.alfi.tmax.NAm.grid.1960.2099 {
dimensions:
    lon = UNLIMITED ; // (462 currently)
    lat = 222 ;
    time = 51100 ;
variables:
    double lon(lon) ;
        lon:units = "degrees_east" ;
        lon:long_name = "lon" ;
    double lat(lat) ;
        lat:units = "degrees_north" ;
        lat:long_name = "lat" ;
    int time(time) ;
        time:units = "days since 1959-12-31" ;
        time:long_name = "time" ;
    double tmax(lon, lat, time) ;
        tmax:units = "degreesC" ;
        tmax:missing_value = 1.e+30 ;
        tmax:long_name = "Downscaled Maximum Temperature in Degrees Celsius" ;
```

THREE. A database of daily maximum and minimum temperature and 24h cumulative precipitation projections for 8,212 (maximum temperature), 8,176 (minimum temperature) and 10,852 (precipitation) individual long-term weather stations in the continental U.S., southern Canada, Alaska, northern Mexico and the Caribbean.

This dataset consists of projections for individual locations, identified by latitude, longitude, and station ID. Although the original input observations are archived in ASCII format, the resulting projections are archived in self-describing netCDF files, one for each variable and climate model/scenario combination. The list of available locations is contained within each self-describing netCDF file and is also provided as an Excel file in **Appendix A**.

Daily historical simulations and future projections were generated for 16 global climate models and 4 future SRES scenarios for the period 1960 to 2099 for a total of 51,100 days per file. The relevant variables in the file are the downscaled projections, contained in tmax/tmin/pr (dimensioned as the number of days by number of stations), and the station information, contained in station_id, latitude, and longitude (dimensioned by the number of stations).

As the files are approximately 1.7 GB each, these are provided in electronic format only.

The format of the files is as follows (sample file header):

```
netcdf pcm.alfi.tmax.NAm.stations.1960.2099 {
dimensions:
    stn_number = 8212 ;
    time = UNLIMITED ; // (51100 currently)
    ncar = 11 ;
    ids = 8212 ;
variables:
    int stn_number(stn_number) ;
        stn_number:units = "" ;
    float longitude(stn_number) ;
        longitude:units = "degrees_east" ;
        longitude:missing_value = 1.e+30f ;
    float latitude(stn_number) ;
        latitude:units = "degrees_north" ;
        latitude:missing_value = 1.e+30f ;
    double time(time) ;
        time:units = "days" ;
    float tmax(time, stn_number) ;
        tmax:units = "see_header" ;
        tmax:missing_value = 1.e+30f ;
    int ncar(ncar) ;
        ncar:units = "" ;
    int ids(ids) ;
        ids:units = "" ;
    char station_id(ids, ncar) ;
        station_id:units = "" ;
}
```

FOUR. Projected changes in seasonal and annual temperature and precipitation and other impact-relevant indicators, per degree global mean temperature change (averaged across scenarios) and for three future time periods compared to a historical time period (for individual scenarios).

One hundred and fourteen seasonal and annual climate indicators were calculated from the gridded downscaled daily maximum and minimum temperature and 24h cumulative precipitation products for the continental U.S., southern Canada and northern Mexico. The indicators include seasonal and annual average temperature and precipitation, as well as secondary indicators such as quantiles, number of days per year above or below a given

threshold, and hybrid indicators such as hot/dry and cool/wet days. The indicators and their units are listed in Table 3.

The original indicators are archived in netCDF format, as follows (sample header):

```
netcdf hadcm3.a2.indicators {
dimensions:
    lon = 460 ;
    lat = 222 ;
    time = UNLIMITED ; // (140 currently)
variables:
    double lon(lon) ;
        lon:standard_name = "longitude" ;
        lon:long_name = "longitude" ;
        lon:units = "degrees_east" ;
        lon:axis = "X" ;
    double lat(lat) ;
        lat:standard_name = "latitude" ;
        lat:long_name = "latitude" ;
        lat:units = "degrees_north" ;
        lat:axis = "Y" ;
    double time(time) ;
        time:standard_name = "time" ;
        time:units = "year as %Y.%f" ;
        time:calendar = "proleptic_gregorian" ;
    float indicator.1(time, lat, lon) ;
        indicator.1:units = "annual value" ;
        indicator.1:_FillValue = 1.e+30f ;
    float indicator.2(time, lat, lon) ;
        indicator.2:units = "annual value" ;
        indicator.2:_FillValue = 1.e+30f ;
    float indicator.3(time, lat, lon) ;
        indicator.3:units = "annual value" ;
        indicator.3:_FillValue = 1.e+30f ;

    ...

    float indicator.114(time, lat, lon) ;
        indicator.114:units = "annual value" ;
        indicator.114:_FillValue = 1.e+30f ;
```

The variables are numbered according to the indicator numbers assigned in Table 3.

Indicator files contain one value for each of 140 years, and one value per grid cell for each of 22 latitude and 460 longitude bands. There are 460 bands instead of 462 because precipitation was downscaled using a 3x3 grid to minimize pixilation during heavy storms (see **ANALYSIS AND FINDINGS** for more information). There is one file for each global climate model/scenario combination. Models with only time slices available (see Table 2) were not included in these calculations as insufficient data was available to calculate the temporal averages.

As these files are approximately 6.5GB each, they are provided in electronic format only.

In addition to providing annual values, projected changes were also averaged for graphical analysis. The absolute values were calculated for the historical period 1960-1979 while simulated changes or anomalies were calculated for one historical and three future climatological 20-year time periods: 1990-2009, 2020-2039, 2050-2069, and 2080-2099.

These averages are provided in netCDF format, one file for each indicator and scenario, as follows:

```
netcdf a2.ts.indicator95 {
dimensions:
    lon = 460 ;
    lat = 222 ;
    time = UNLIMITED ; // (5 currently)
variables:
    double lon(lon) ;
        lon:units = "degrees_east" ;
    double lat(lat) ;
        lat:units = "degrees_north" ;
    double time(time) ;
        time:units = "years" ;
    float mod0(time, lat, lon) ;
        mod0:units = "annual value" ;
        mod0:missing_value = 1.e+30f ;
    float mod1(time, lat, lon) ;
        mod1:units = "annual value" ;
        mod1:missing_value = 1.e+30f ;
    float mod2(time, lat, lon) ;
        mod2:units = "annual value" ;
        mod2:missing_value = 1.e+30f ;

    ...
    float mod17(time, lat, lon) ;
        mod17:units = "annual value" ;
        mod17:missing_value = 1.e+30f ;
}
```

where “mod0” refers to the all-model average and the remaining models are listed in alphabetical order as they appear in Table 2. The files are named according to the indicator numbers assigned in Table 3.

The majority of the indicator files contain values for BCC-BCM2 (mod1 out of mod17); however, it is not recommended that these be used, as the BCC-BCM2 model showed consistently large biases in simulating historical temperature and precipitation and was ultimately excluded from the dataset on this basis (see **ANALYSIS AND FINDINGS** for more information).

The results of the temporal averaging are summarized in a series of regional temperature-precipitation scatter plots (**Appendix B**; original Excel files also provided). These plots show the spatially-averaged projected change in temperature and precipitation for the three future time periods (2020-2039, 2050-2069 and 2080-2099), compared to the historical 1960-1979 period, for 8 U.S. regions: Pacific Northwest, Southwest, Northern Great Plains, Southern Great Plains, Midwest, U.S. Northeast and Canadian Maritimes, Mid-Atlantic, and Southeast. In addition to the scatter plots for temperature and precipitation, each Excel file also contains a map defining the geographic boundaries of each region, and a scatter plot of the projected change in the number of hot/dry and cool/wet days per year.

Table 3. Primary and secondary climate indicators calculated for 5 twenty-year time slices (1960-1979, 1990-2009, 2020-2039, 2050-2069, and 2080-2099) and 4 twenty-year global mean temperature targets (+0,1,2,3°C relative to 1971-2000) and their units. Indicator numbers refer to the title of the netCDF files.

Indicator	Units
ANNUAL AND SEASONAL AVERAGES	
1-5. Average maximum temperature (winter or DJF; spring or MAM; summer or JJA; fall or SON; annual)	degrees C
6-10. Same, for average minimum temperature	"
11-15. Same, for cumulative precipitation	mm per season or year
ANNUAL AND SEASONAL EXTREMES	
16. Hot summers: Average maximum temperature exceeding hottest historical (1971-2000) summer	Number of years historical threshold exceeded (maximum 20)
17. Cold years: Average minimum temperature below coldest historical (1971-2000) year	"
18-19. Wet/dry years: Cumulative precipitation above/below historical (1971-2000) wettest/driest year	"
ANNUAL AND SEASONAL VARIANCE	
20-24. Standard deviation of daily maximum temperature (winter or DJF; spring or MAM; summer or JJA; fall or SON; annual)	degrees C
25-29. Same, for daily minimum temperature	"
30-34. Same, for cumulative precipitation, wet days only	log(mm)
35-39. Range in daily maximum temperature (winter or DJF; spring or MAM; summer or JJA; fall or SON; annual)	degrees C
40-44. Same, for daily minimum temperature	"
45-49. Precipitation intensity, defined as cumulative precipitation divided by number of wet days (winter or DJF; spring or MAM; summer or JJA; fall or SON; annual)	mm/day
ANNUAL EXTREMES	
50-53. Temperature of 1, 3, 5, and 10-day hottest period (based on daily maximum temperature)	degrees C
54-57. Same, for coldest period (based on daily minimum temperature)	"
58-61. Same, for cumulative precipitation	mm
62-66. Total number of dry days defined as $pr < 0.01$ " (winter or DJF; spring or MAM; summer or JJA; fall or SON; annual)	Average number of days per year
67-71. Longest stretch or period of dry days, defined as $pr < 0.01$ " per day (winter or DJF; spring or MAM; summer or JJA; fall or SON; annual)	Total number of days per year
72-76. Average length of dry periods, defined as $pr < 0.01$ " per day (winter or DJF; spring or MAM; summer or JJA; fall or SON; annual)	"
77-79. Beginning, end, and length of growing season (defined as the period between the last frost in spring and the first frost in autumn)	Julian day (for beginning and end of season) and number of days (for length of season)
THRESHOLDS	
80-84. Days per year above 32, 65, 80, 90, and 100°F, based on daily maximum temperature	Average number of days per year
85-86. Days per year below 20, 32°F, based on daily (nighttime) minimum temperature	"
87-89. Nights per year above 65, 75, 85°F, based on daily (nighttime) minimum temperature	"
90. Wet days ($pr > 0.01$ ")	"
91-93. Heavy precipitation days ($pr > 1, 2, 3$ inches)	"

94-95. Cumulative growing degree-days, based on average temperature ((max+min)/2), using a threshold of 65, 75°F	Degree-Days
QUANTILES	
96-100. Number of hot days per year with daily maximum temperature above the historical (1971-2000) 50, 75, 90, 99 and 99.9 th quantile	Average number of days per year
101-105. Number of cold nights per year with daily minimum temperature below the historical (1971-2000) 0.1, 1, 10, 25, and 50th quantile	"
106-110. Number of wet days with precipitation exceeding the historical (1971-2000) 50, 75, 90, 99 and 99.9 th quantiles	"
HYBRIDS	
111. Number of snow days per year (defined as daily maximum temperature equal to or below freezing and precipitation above 0.01")	Average number of days per year
112. Ratio of precipitation falling as rain to that falling as snow (defined as daily maximum temperature equal to or below freezing)	Ratio (rain/snow), mm/mm
113. Number of hot, dry days per year (defined as precipitation < 0.01" and daily maximum temperature > 90°F)	Average number of days per year
114. Number of cool, wet days per year (defined as precipitation > 0.01" and daily maximum temperature < 65°F)	"

Using the approach introduced by the 2011 National Research Council Report, *Climate Stabilization Targets: Emissions, Concentrations, and Impacts over Decades to Millennia* (NAS, 2011), future projections were also averaged **across scenarios** by calculating the changes expected for global mean temperature changes of +1°C, +2°C and +3°C compared to historical 1971-2000 for each individual simulation. This approach has the advantage of being able to compare projected changes across a range of scenarios and climate model sensitivities to resolve the differences in the magnitude and pattern of expected change that is *independent* of the uncertainty in either human scenarios or climate sensitivity. This approach also presents impacts within a policy-relevant framework, as national and international agreements (such as the EU target of +2°C) can be couched in terms of global mean temperature targets rather than in terms of a given emissions scenario.

Calculating projected changes under a series of global mean temperature targets was accomplished as follows. First, annual average global average temperature was calculated for each model for each year from 1960 to 2099. Then, the 20-year running mean values were calculated for each year beginning with 1960-1979 and ending with 2080-2099. Next, the first period in which the 20-year global mean temperature was equal to or higher than +1/2/3°C compared to the 1971-2000 global mean temperature *of that simulation* was identified (see Table 4). Finally, this date was used to select the 20-year period from the corresponding simulation to average for each indicator.

As for the time averages, these global mean temperature values are provided in netCDF format, one file for each indicator, with each indicator having a "time" dimension of 4, corresponding to +0 (absolute value) and +1, 2, and 3°C (anomaly relative to simulated 1971-2000). The files are named according to the indicator numbers assigned in Table 3. Again, models with only time slices available (see Table 2) were not included in these calculations as insufficient data was available. The results are also presented in graphical format in **Appendix C** and copies of the original postscript plots are provided electronically.

Table 4. First year in which the twenty-year climatological average global mean temperature for each global climate model/scenario simulation beginning with that year is equal to or more than +1, 2, and 3°C above the 1971-2000 average for that simulation. Earlier years indicate a more sensitive climate model, later years a less sensitive one. For the lower emissions scenarios (e.g. A1B, B1) some models do not reach +2 or +3°C before end-of-century. These are indicated by an “NA” value.

	+1oC	+2oC	+3oC		+1oC	+2oC	+3oC
SRES A1FI				SRES A1B			
ccsm	2009	2036	2054	bcm2	2031	2056	NA
gfdl_2.1	2014	2040	2061	ccsm	2010	2043	NA
hadcm3	2023	2044	2060	cgcm3_t47	2018	2057	NA
pcm	2023	2052	2077	cgcm3_t63	2010	2042	2075
SRES A2				cnrm	2020	2052	NA
bcm2	2030	2059	2080	echam5	2026	2051	2073
ccsm	2009	2043	2068	echo	2025	2053	NA
cgcm3_t47	2020	2052	2075	gfdl_2.1	2022	2054	NA
cgcm3_t63	2008	2040	2062	hadcm3	2022	2052	NA
cnrm	2027	2051	2070	hadgem	2017	2045	2070
echam5	2030	2054	2072	miroc_med	2016	2046	2074
echo	2027	2054	2078	pcm	2028	NA	NA
gfdl_2.0	2021	2051	2073	SRES B1			
gfdl_2.1	2025	2057	2080	bcm2	2039	NA	NA
hadcm3	2026	2051	2074	ccsm	2014	NA	NA
hadgem	2018	2045	2066	cgcm3_t47	2027	NA	NA
miroc_med	2015	2050	2068	cgcm3_t63	2011	2061	NA
pcm	2033	2068	NA	cnrm	2029	NA	NA
				echam5	2035	2069	NA
				echo	2033	NA	NA
				gfdl_2.0	2021	2079	NA
				gfdl_2.1	2026	NA	NA
				hadcm3	2036	NA	NA
				miroc_med	2022	2068	NA
				pcm	1995	NA	NA

ANALYSIS AND FINDINGS

The previous section described the specific products generated by this project. This section focuses on the lessons learned while creating and analyzing these products, many of which are generally applicable to any generation of high-resolution climate projections for impact assessments. This section also discusses the innovative approaches taken that expanded the original project scope in order to address questions and concerns that arose during the course of the research.

The lessons learned and solutions achieved accomplished can be broadly categorized under five topics:

1. Global climate model and observational inputs: quality control and preparation to improve quality of downscaled output
2. Downscaling method: modifications to improve quality of output
3. High-resolution climate projections: comprehensive error and bias tests to quantify quality of output
4. Computation: technical improvements to methods and code to increase efficiency
5. Education: two-way education of the users--and the downscalers—on the application of this dataset to impact analyses

Each of these issues and the solutions that were identified and implemented are discussed in detail below.

Improvements to Downscaling Inputs

Quality Control of Global Climate Model Output

Global climate model outputs used as predictors for downscaling were carefully reviewed and quality-controlled to identify errors and peculiarities in the output fields that would need to be documented, corrected, and/or standardized before use in downscaling.

The majority of global climate model fields used in this analysis were obtained from the CMIP3 archive, which maintains standards of data provision and quality control. (SRES A1FI scenarios were not archived by CMIP3: these simulations were obtained directly from individual modeling groups.) In the past, it has been taken for granted that, because of the standards required for modeling output to be submitted to CMIP3, these fields can be used with confidence. This analysis, however, has shown that such an assumption is not valid.

Before the global model fields could be used, they had to be carefully quality-controlled for a number of specific problems that were identified by trial and error. Specific problems identified in the course of this analysis consisted of:

1. Missing data (which could range from a few random days to entire decades in the middle of a simulation)
2. Incorrect values (unrealistic data points far above or below historical observed maxima or minima)
3. Mis-labeled files (files that did not contain the variable or data listed in the file name and header)

This last issue was particularly insidious and difficult to identify, requiring communication with the original modeling group in order to confirm the existence of incorrectly-labeled files.

In addition to these errors, each model also had its own peculiarities that had to be standardized before the predictors could be incorporated into the downscaling model. For example, most models have 365 days per year with no leap year. A few models did have a leap year, while other models only have 360 days per year. All models had to be standardized to 365 days per year, in some cases by removing the leap year days and in others, by inserting “NA” values for the 31st of each month from August to December. Most models archived data running from south to north and east to west; however, some model outputs ran from north to south. Some models archived precipitation in m/s, others in kg/m²/s. Five models only had limited future time slices available, others began in 1961 instead of 1960.

These and other known quality issues and model peculiarities are listed in Table 5.

Table 5. Known quality control problems with original global climate model outputs and model peculiarities that had to be documented and, if possible, standardized before the model could be used in downscaling.

Global climate model	Known issues
BCM2	A1fi scenario not available. Leap years removed before downscaling. Erroneous data points in file above maximum observed values. Model performance significantly poorer than others; removed from final data provision.
CCSM3	Temperature data for historic and future periods not initialized from the same run, creating disconnect between the two files if one were to plot a continuous time series joining the 2 files. This situation is unavoidable for CCSM, as most of its historical simulations have erroneous tmax/tmin values (they were accidentally overwritten with the variable TREFHT by the original modeling group). A1fi Tx/Tn: b30.030h to b30.099a (no match) A1fi Pr: b30.030a to b30.099a (matched) A1B Tx/Tn: b30.030h to b30.042g (no match) A1B Pr: b30.040a to b30.030a (matched) A2 Tx/Tn: b30.030h to b30.042e (no match) A2 Pr: b30.030e to b30.042e (matched)
CGCM3-T47	A1fi scenario not available. Simulation begins in 1961. 1960 is missing.
CGCM3-T63	A1fi scenario not available. Simulation begins in 1961. 1960 is missing.
CNRM	A1fi scenario not available. Leap years removed before downscaling.
CSIRO	A1fi scenario not available. Future data not available for 2000-2045 and 2066-2080.
ECHAM5	A1fi scenario not available. Leap years removed before downscaling.

ECHO	A1fi scenario not available. Model has 360 days per year. These were divided up such that the last days of May, July, Aug, Oct and Dec are always missing .
GISS-AOM	A1fi and A2 scenarios not available. Future data not available for 2000-2045 and 2066-2080.
GFDL CM2.0	A1fi and A1B scenarios not available.
GFDL CM2.1	No known issues.
HadCM3	Model has 360 days per year. These were divided up such that the last days of May, July, Aug, Oct and Dec are always missing . For the A1B simulation, the last 30 days are missing (Dec 2099) and 10 years between 2080-2089 are missing for precipitation. For A2 and B1, the last 60 days of the historical period (Nov-Dec 1999) are missing and the first 10 days of Nov 2000 are missing. For B1, the year 2000 is missing.
HadGEM	Model has 360 days per year. These were divided up such that the last days of May, July, Aug, Oct and Dec are always missing . A1fi and B1 scenarios not available.
Miroc-med	A1fi scenarios not available. Future data not available for 2000-2045 and 2066-2080. Leap years removed before downscaling.
Miroc-hi	A1fi and A2 scenarios not available. Future data not available for 2000-2045 and 2066-2080. Leap years removed before downscaling.
MRI_CGCM2	A1fi scenario not available. Future data not available for 2000-2045 and 2066-2080.
PCM	PCM A2, A1B and B1 scenarios are based on B06.08 to 1980 and B07.08 from 1980 to 2099 but daily data for B06.08 is not available before 1980 so B06.57 was used instead. There will be a level shift at 1980 because of this. The last 61 days of 2099 (Nov-Dec) are also missing in the A1B scenario.

Filtering for Low-Level Noise in Temperature Fields

A new pre-processing approach was developed to remove spurious noise in global climate model daily temperature fields. This approach can be used by any downscaling method in the future. Its application significantly improved biases in downscaled temperature, particularly for cold and hot extremes.

It is well-known in the climate community that ensemble mean GCM projections typically out-perform any individual model simulation (Knutti et al., 2010). One of the reasons this may be the case is because the ensemble mean represents the smoothed average of multiple global climate model fields, and others have recommended smoothing GCM output before using it to assess regional climate change (e.g., Raisanen and Ylhaisi, 2011). For this reason, as part of this project we developed and tested a new pre-processing approach for daily GCM temperature fields to determine whether removal of low-level and potentially spurious noise from global climate model outputs would improve the quality of the resulting high-resolution projections.

We first developed a flexible code where the variance of GCM temperature fields can be deconstructed using Empirical Orthogonal Function (EOF) analysis for each of the synoptic-scale regional boxes shown in Figure 8 (Alaska, continental U.S., Caribbean). Synoptic scale is defined by NOAA as “the spatial scale of the migratory high and low pressure systems of the lower troposphere, with wavelengths of 1000 to 2500 km”. The primary function of the GCMs, as used in this analysis, is to simulate the effects of climate change on these large-scale weather patterns; hence the need to analyze model variance *at the skillful scale of the output fields*.

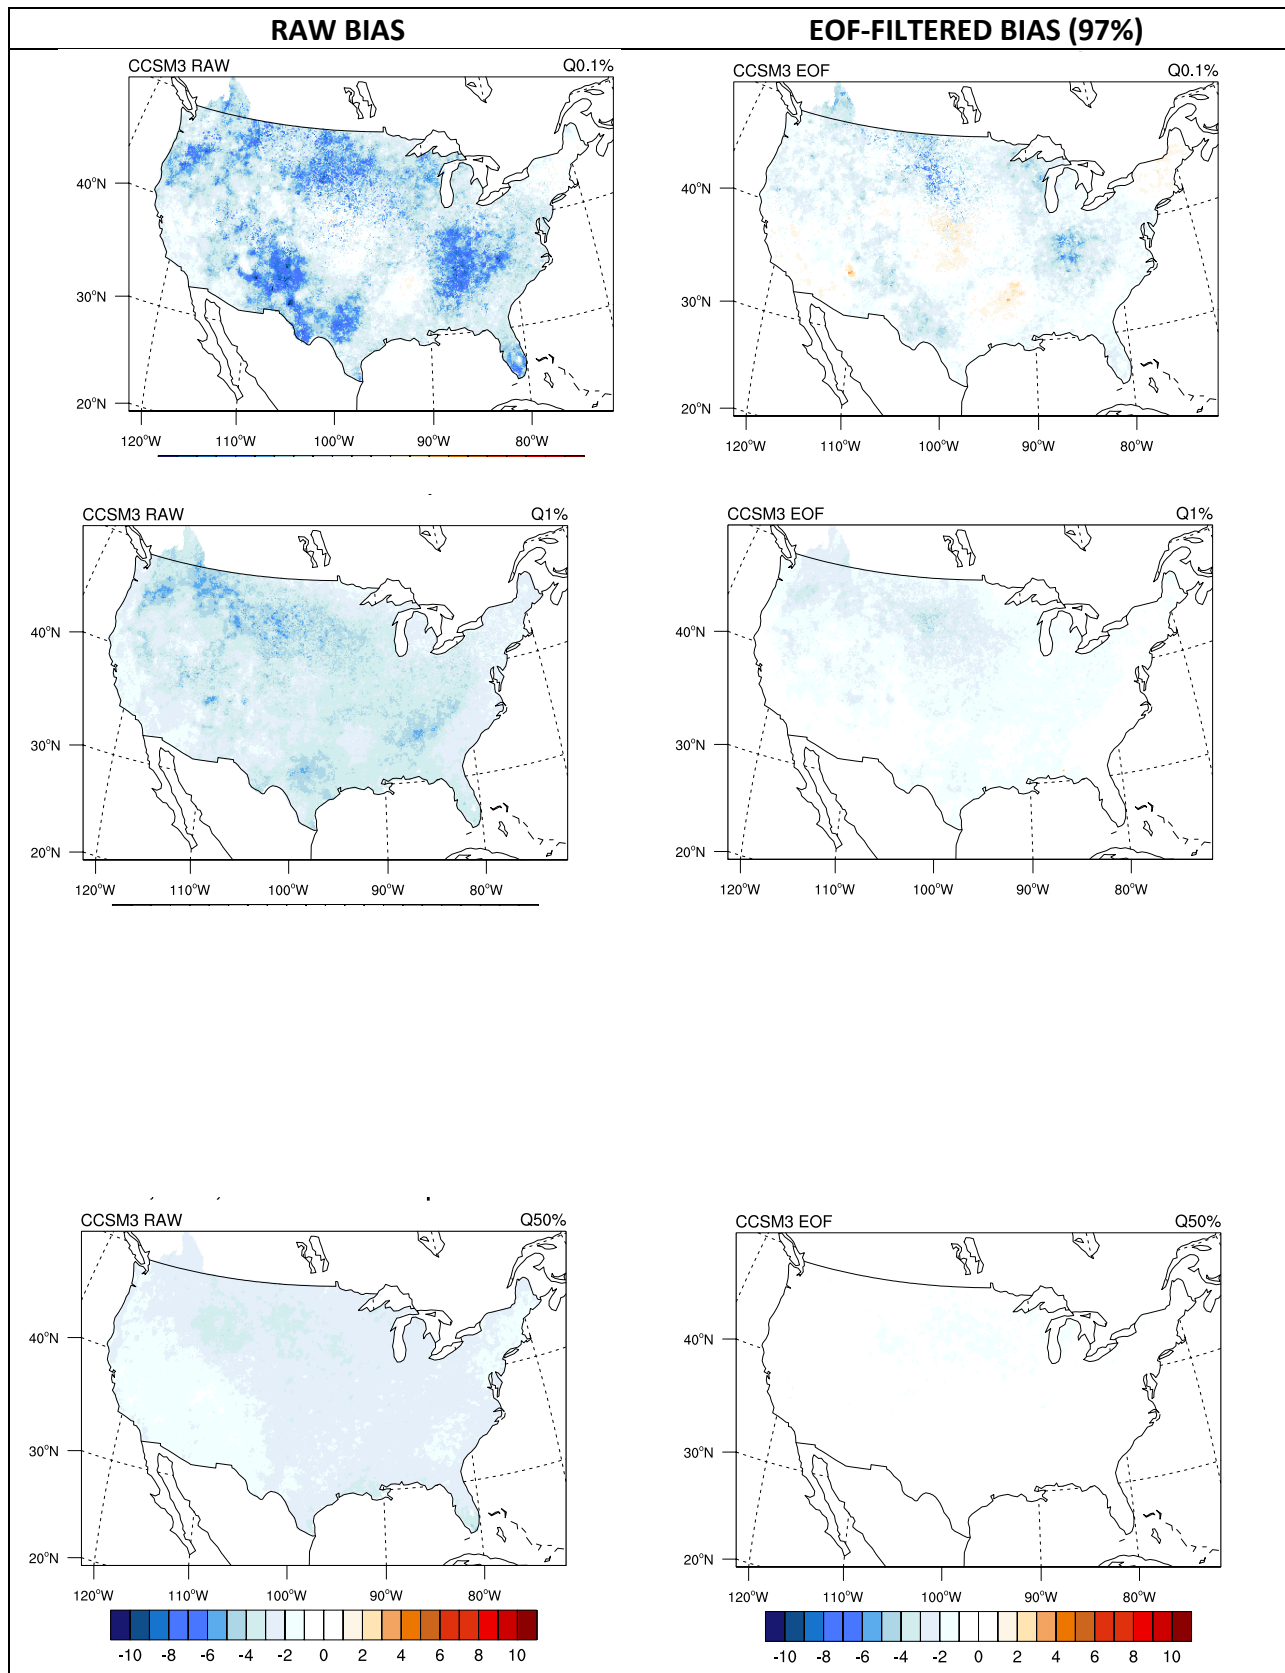
Part of the simulated variance was removed before reconstructing the original temperature fields for each of these large or synoptic-scale regions from a smaller sub-set of principal components than were derived from the original data. These reconstructed temperature fields were then used to downscale both individual station locations as well as the entire continental U.S. high-resolution grid for a limited subset of global model inputs. Through iterating between a range of 95% to 99% of retained variance, it was determined through root mean square errors (RMSEs) and mean absolute error (MAE) that 97% was a generally appropriate threshold that minimized error across the continental U.S., with both higher and lower thresholds resulting in overall higher errors.

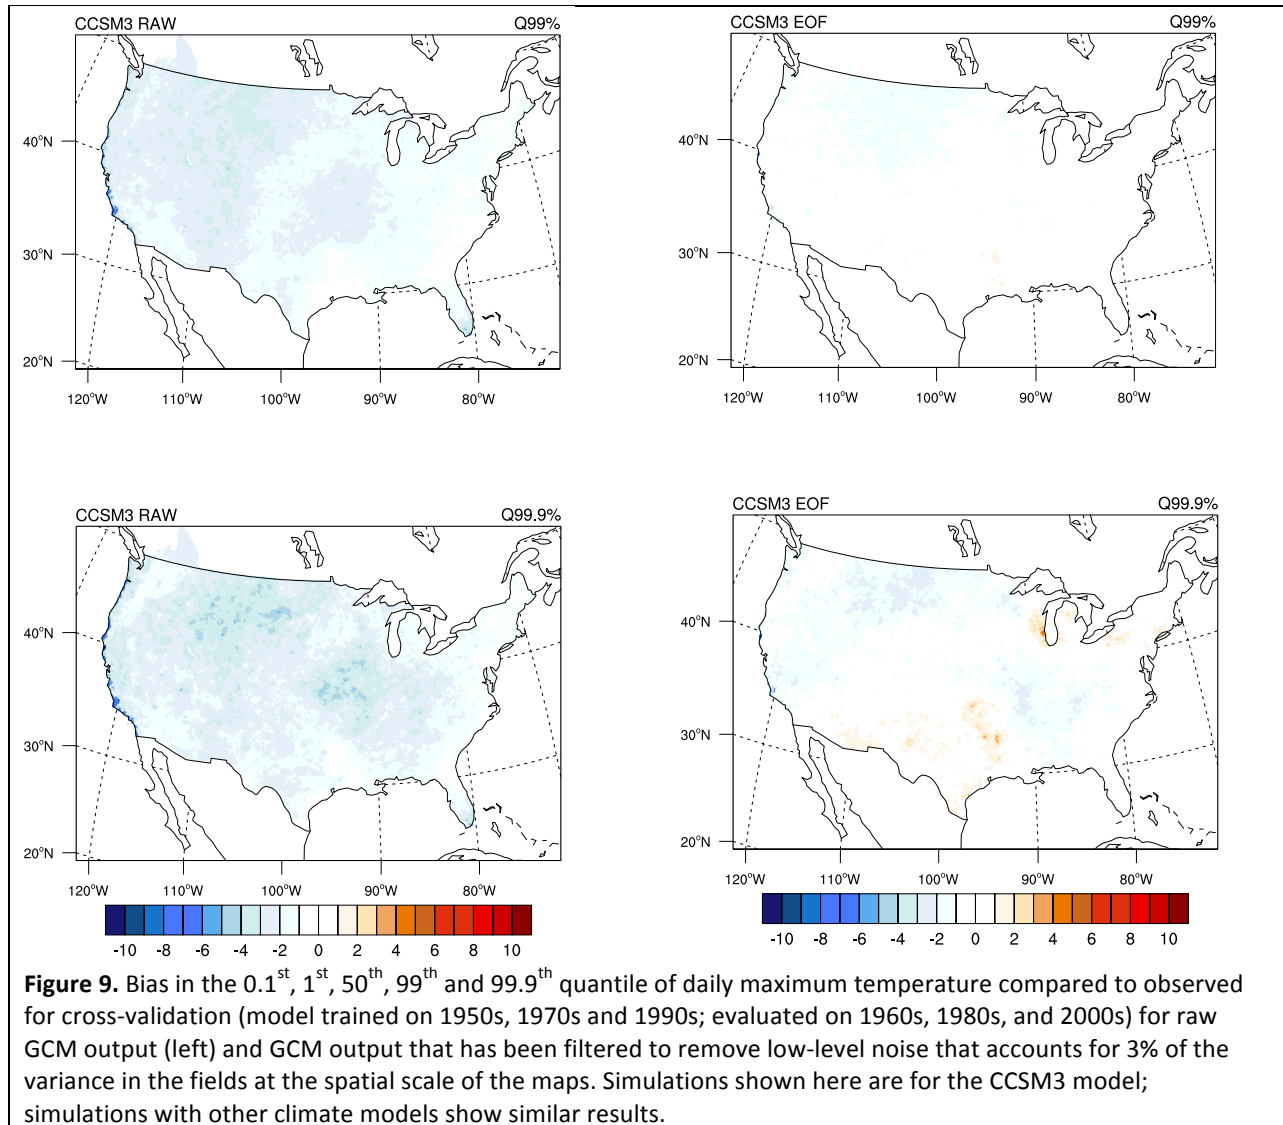
This sensitivity analysis revealed a number of interesting and unique results. First and most importantly, the overall effect of this filtering was to improve the quality of the temperature downscaling. Figure 9 compares the bias in simulated vs. observed maximum daily temperature from quantiles ranging from 0.1 to 99.9. In both cases, the downscaling model was trained on odd decades (1950s, 1970s and 1990s) and evaluated relative to independent observations in even decades (1960s, 1980s, 1990s). For the unfiltered (raw) data, there are significant and wide-spread biases at the tails of the distribution. These are noticeably reduced when the global climate model output is filtered before downscaling.

It is important to note that the improvement varied by location, with the most significant improvements being for inland locations with higher variance in the daily temperature record. Improvement also varied over the distribution, with the greatest improvement due to EOF-filtering (but still greatest remaining bias) being seen at the tails of the distribution as compared to the mean; and by climate model, with some showing greater improvements than others.

For the station downscaling, we also found the largest improvement was achieved for the stations with the shortest and poorest-quality records. By reducing spurious noise in the global climate model fields, this filtering process enables the statistical downscaling model to better utilize limited observational data and improve model training during the historical period.

We are continuing to apply this EOF filtering method to CMIP5 temperature output fields. **To our knowledge, this innovative filtering method is not yet in use by any commonly used downscaling method in the U.S.; however, it could be easily incorporated into any downscaling framework.**





Compared to temperature, precipitation tends to display a greater amount of smaller-scale variability in both observations and GCM output fields. The EOF filtering approach was also tested on precipitation, but in this case even a low level of noise removal was found to degrade rather than assist precipitation downscaling. For that reason, no filtering was conducted on the precipitation input fields

Spatial Interpolation of Predictor Fields

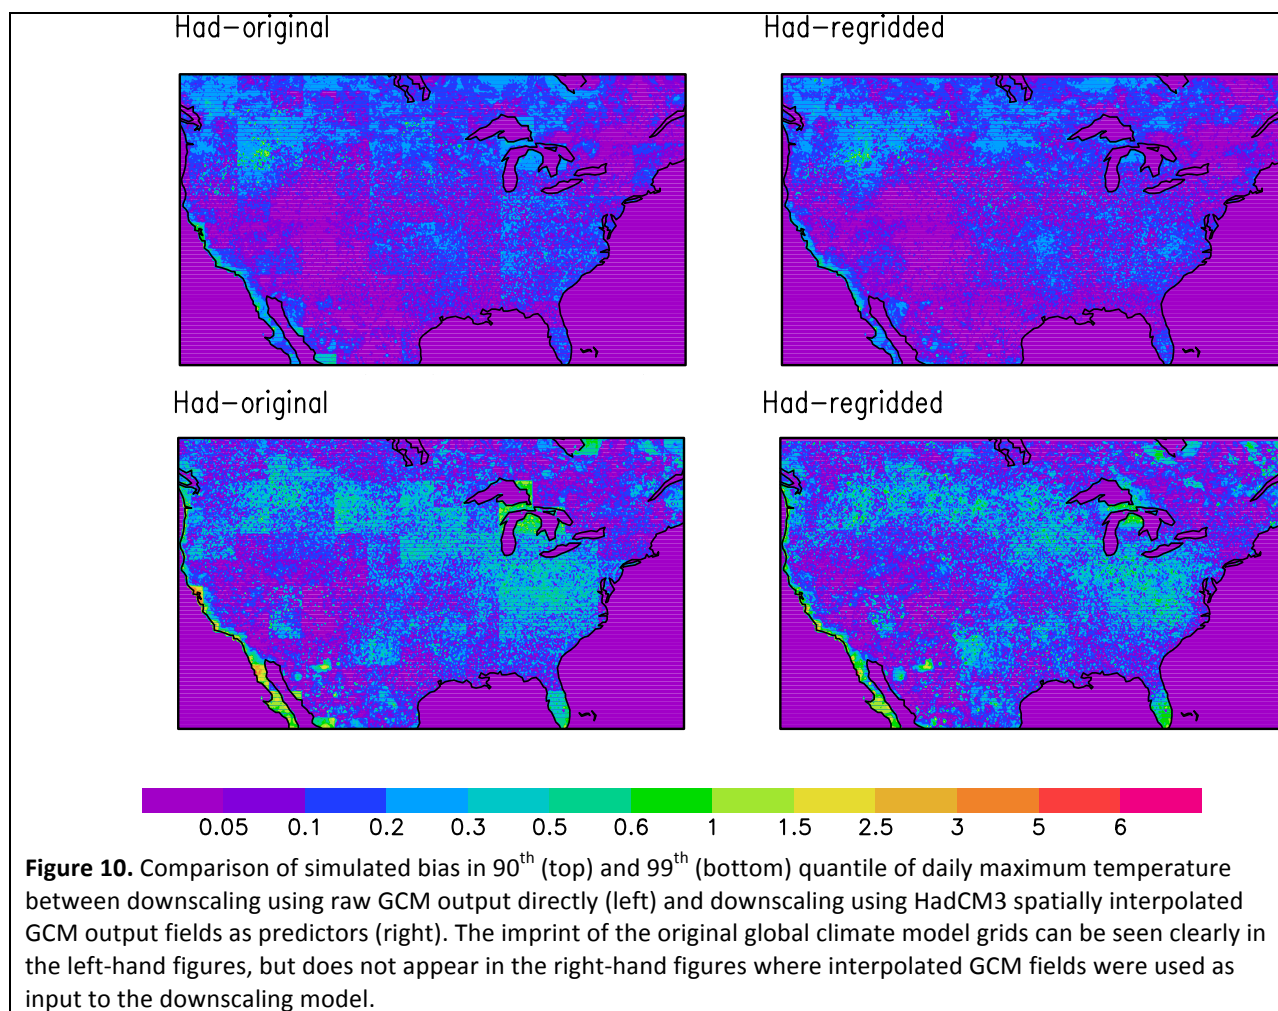
Spatial interpolation of raw global climate model output fields to the scale of the high-resolution gridded observations significantly improved downscaling performance by creating a smooth transition in the shape of the GCM quantile-OBS quantile relationship from one grid cell to the next.

Even after filtering for both temperature and precipitation, initial downscaling of GCM fields showed that an imprint of the original GCM grid cell often remained. This imprint was not generally visible in annual or seasonal mean values, but rather in thresholds (e.g. number of wet days per year) or extremes (e.g. number of days per year over 90°F). The

reason for this remaining “imprint” was determined to be the result of a sharp change in the **shape** of the distribution of daily temperature and precipitation from one GCM grid cell to the next. With a quantile regression-based downscaling approach, smooth changes in the shape of the probability distribution over space are required to produce smooth output fields.

To address this issue, specifically to improve the homogeneity of the projections towards the tails of the distributions (e.g. cold/hot or wet/dry), the original downscaling was completely re-done after all global model output fields were linearly interpolated to the scale of the one eighth-degree grid before being used as input to the downscaling routine. These interpolated fields were also used as input to station-based downscaling. Figure 10 provides a side-by-side comparison of downscaled projections before and after the spatial interpolation of the predictor fields.

*As with the EOF filtering method, this simple interpolation of global climate model-based predictors can be used by any downscaling method to improve the spatial homogeneity of their results. **Its application significantly improved model performance, particularly for cold temperatures and for spatial patterns of precipitation.***



Quality Control of Weather Station Observations

A rigorous quality control and nearest-neighbor process was developed to identify and remove erroneous observed temperature and precipitation values. This is a generalizable code that can be used to quality control any observational dataset.

Most government agencies perform some operational QC on observational data before it is released. In terms of the data sets used here, the Global Historical Climatology Network (GHCN; <http://www.ncdc.noaa.gov/oa/climate/ghcn-daily/>) contains daily maximum and minimum temperature and daily cumulative precipitation. According to their metadata, GHCN station data undergoes extensive testing for operational QC including checks for:

- Duplicate data
- World record exceedence
- Consecutive identical values
- Frequent-value check (precipitation only),
- Gap check, Z-score-based climatological outlier check,
- Internal consistency
- Temporal consistency (spike or dip),
- Lagged temperature range
- Spatial consistency (corroboration of anomalies)

The National Climate Data Center (<http://www.ncdc.noaa.gov/ghcnm/v3.php>) quality control tests include:

- Consecutive month duplication
- Series duplication
- Streak
- Isolated value
- Climatological outlier
- Spatial inconsistency


The U.K. Meteorological Office Integrated Data Archive System land surface observational database (MIDAS, http://badc.nerc.ac.uk/view/badc.nerc.ac.uk_ATOM_dataent_ukmo-midas/), contains daily mean temperature and daily cumulative precipitation. Midas data also includes raw temperature readings taken every 3 hours or every 6 hours depending on the observation location and the year. Midas data, like GHCN and NCDC, is checked extensively prior to release (http://badc.nerc.ac.uk/data/surface/ukmo_guide.html#7).

Midas tests include:

- Ingestion checks
- Checks against neighbors
- Manual quality control
- Gross errors

Despite the fact that most of the station data used in this project had undergone a standardized quality control, during the course of the project there were numerous instances of clearly erroneous values (e.g. very cold days in warm locations, days where maximum temperature was lower than minimum temperature, days with negative precipitation) in the observations that strongly degraded the quality of the statistical downscaling model.

The disproportionate impact of even a single erroneous daily value on a quantile regression-based statistical model can be illustrated by this example from southern Florida. The GHCN location of Hialeah, a suburb of Miami, had one data point with daily temperature 20°C colder than any other day at that location or any nearby location. When that data point was included in training the full ARRM model (indicated by “full” in Figure 11), the RMSE in the resulting cross-validated downscaling was nearly an order of magnitude higher (2°C instead of ~0.2°C, left figure) than when that single data point was removed (right figure).



Prior to commencing the project, it was assumed that quality-controlled station data would be sufficient; as with the global climate model outputs, that assumption was clearly not valid and had to be addressed directly before station data could be reliably downscaled.

To address, at minimum, the most readily-identifiable problems with individual daily values in any observational dataset, we greatly expanded the original scope of this project to develop our own comprehensive weather station data quality control process. We also obtained and incorporated additional datasets (NCDC, MIDAS, etc.) into the database of station observations to provide maximum density for a nearest-neighbor test of outlier values at any individual station.

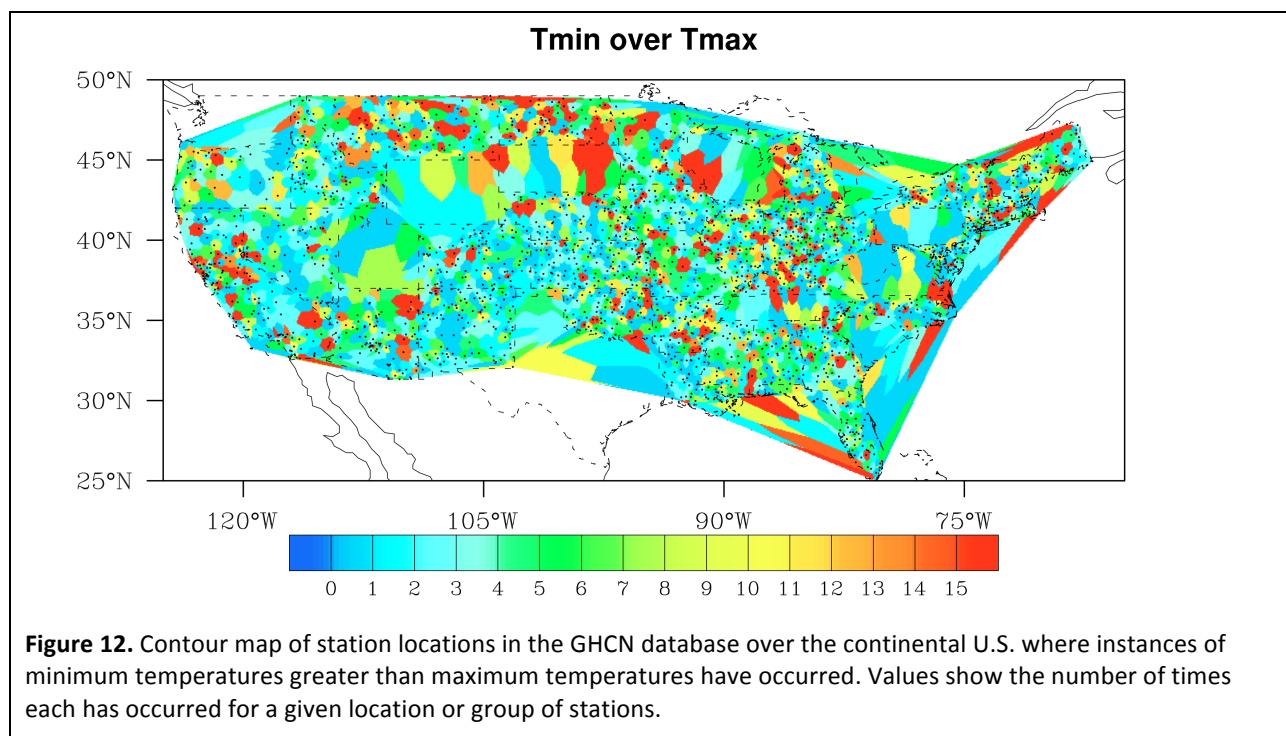
The first step in our quality control process was to create a master file documenting the name, latitude, and longitude of every station in the various datasets available. Using latitude and longitude values, stations within 1km of each other were assumed to be co-located and only the longest unique station record for each variable was retained.

The quality control process consists of two steps: first, individual quality control for each station; and second, a nearest neighbor approach to validate outliers identified relative to the climatology of each month. Individual quality control identified and replaced with N/A any values that failed the following tests:

- **Tmin>Tmax** – Daily minimum temperature exceeds the daily maximum temperature.
- **Realistic Range** – Daily reported minimum temperature exceeds the reported maximum, any temperature values above (below) the highest (lowest) recorded values for North America (-50 to 70°C) or with precipitation below zero or above the highest recorded value for the continental U.S. (915 mm in 24h).
- **Repeated Values** – Daily maximum temperature, daily minimum temperature, daily average temperature, or non-zero daily precipitation values repeat for 5 or more consecutive days to within one tenth of a mm per day or one hundredth of a degree Celsius.

Unlike many of the preceding research efforts into data quality control, this process removes data when errors are found without replacing the data with new values and is fully automated, requiring no labor-intensive manual verifications or judgments.

Many of these tests appear redundant with operational QC already performed on these datasets. However, preliminary testing revealed Tmin>Tmax examples in every dataset, even those that explicitly test for it (Figure 12). Realistic Range errors are also common in all datasets. Overall, this automated, standardized quality control process flagged errors in 99% of the GHCN observation locations and 10% of the MIDAS observation locations tested and removed 0.4% of the GHCN temperature data and 0.01% of the MIDAS temperature data.



In the second step of the quality control process, the merged database of station locations is first used to identify up to 10 “nearest neighbors” for each individual weather station within 100km of its location. For each weather station, the monthly (for temperature) and seasonal (for precipitation) distributions are ranked and the highest and lowest N values are identified. The nearest neighbor stations are then queried to see if the days on which values 1 through N occur are also days in which the highest M values for the neighbor station occur, plus or minus one day on either side to account for weather systems which may be moving through the area close to midnight. Here, N is taken to be 100 and M to be 500. If any value of N does not occur within +/- one day of a value M for any neighboring station, the value of N is replaced with an NA.

For the North-Central America region shown in Figure 8, the nearest neighbor test identified 82,515 non-corroborated outlier values for maximum temperature for a removal rate of 0.0386% and 92,437 non-corroborated outlier values for minimum temperature for a removal rate of 0.0436%. All outlier values for precipitation were corroborated by neighboring stations.

While comprehensive, this quality control process can not correct for all errors in the data. One error that was identified but not corrected are the spikes of values at 0 and 32 degrees that may be introduced by repeated (and faulty) units conversion. Temperature data can be in Fahrenheit or Celsius and in tenths of degrees or degrees. Precipitation data can be in tenths of inches or millimeters. For consistency within this dataset, all data was converted to tenths of degrees Celsius for temperature and mm/day for precipitation. However, much of the data had likely already undergone multiple conversions (particularly that of the MIDAS database) and a number of stations already displayed a “binned” effect indicating loss of information in transcription and/or translation, as well as the possible introduction of erroneous values such as suspicious spikes at 0 and 32°C (Figure 11).




Figure 11. Histogram of raw daily maximum temperature values from the MIDAS dataset, in degrees Celsius, show clear bins where information has likely been lost in units conversion and potential errors may have been introduced relative to 0°C and 32°F (note spike at 32°C in left-hand figure and 0°C in right-hand figure).

The quality-controlled daily temperature and precipitation have now been downscaled and are incorporated into the database of high-resolution projections, expanding the original geographic scope of the analysis to encompass the entire continent of North America. This dataset has not been made available for public distribution as part of this project since it contains largely original data that requires permission from the original agency, but it has been used as input to the station-based downscaling that is available for distribution as a product of this work.

Improvements to Downscaling Methodology

New predictors for precipitation

The downscaling model now allows multiple precipitation predictors to improve simulation of seasonal precipitation.

The ARRM downscaling model was modified to allow for multiple precipitation predictors (convective, large-scale, and total precipitation). This modification enabled the model to better resolve the drivers of precipitation in summer vs. winter and in warmer vs. colder climates, and reduced biases in both seasonal and extreme precipitation. Models with multiple precipitation predictors consisted of: CCSM3, CNRM, ECHAM5, ECHO, HadGEM and PCM. For the remaining models, the original predictor of total precipitation continued to be used.

3x3 downscaling of precipitation

Initial precipitation downscaling showed unacceptable levels of pixilation in the resulting daily precipitation fields for high precipitation events. This problem was solved by incorporating a 3x3 matrix into the downscaling and an analysis code developed to quantify the resulting improvement in spatial autocorrelation of daily precipitation fields.

In the original dataset, individual rainy days with high precipitation over large spatial areas showed unacceptable levels of pixilation, with adjacent grid cells showing both higher and lower precipitation amounts than a similar event in the historical record. This problem was the result of the downscaling approach treating every grid cell as an individual data point under the assumption that the global climate model output fields were able to provide the spatial coherence and homogeneity in the resulting downscaled high-resolution projections. Although this assumption held true for temperature, it clearly failed for high-precipitation events at the tail of the distribution, where small differences in historical data points (possibly only one or two) could produce large differences in total precipitation for adjacent grid cells.

To solve this problem, the ARRM downscaling model was modified to train gridded precipitation models on a 3x3 grid in order to stabilize the statistical fit at the tails of the distribution. All gridded precipitation downscaling was re-run from scratch with the new code. The modification to the code significantly increased the run time of the statistical model due to the greater volume of data used as input to train the statistical model, but was successful in creating a much more stable and homogeneous fields of daily precipitation.

The problem of pixilation is common in the field of image processing, where it is typically referred to as “salt and pepper”. **Moran’s I** is a measure of spatial autocorrelation that is often used to assess the pixilation of an image by assessing the differences between a grid of pixels. This statistical measure is defined as:

$$I = \frac{N}{\sum_i \sum_j w_{ij}} \frac{\sum_i \sum_j w_{ij} (X_i - \bar{X})(X_j - \bar{X})}{\sum_i (X_i - \bar{X})^2}$$

where N is the number of spatial units indexed by i and j ; X is the variable of interest; \bar{X} is the mean of X ; and w_{ij} is an element of a matrix of spatial weights (Source: Wikipedia).

To quantify the extent to which the 3x3 downscaling approach actually reduced pixilation and the extent to which model output remained biased relative to observations, an analysis code calculating the spatial dependence of daily precipitation fields in terms of the visual appearance of rain storms, the distance between pixel values, and Moran's I, was developed. *This is a generalizable analysis package that can be used to test and compare the coherence of any spatial data field.*

The analysis code first identifies the top 100 wettest rain events in the continental U.S. in each historical simulated and observed period (1960 to 1999) and cuts out a regular square of 81x81 one-eighth degree grid cells centered on each storm. Global climate models are initialized in the late 1800s and thereafter permitted to develop their own patterns of natural variability. Although models are expected to produce climatological conditions (averaged over multiple decades) that resemble observed, they are not and should not be expected to reproduce any specific rainstorms or other weather events on the actual day they occurred. For that reason, rain events must be identified individually for each model simulation.

The analysis code then maps individual events, so they can be compared visually; calculates the differences between individual grid cells and plots a histogram of those differences for both the top10 and top 100 wettest events; and finally calculates the value of Moran's I for each gridcell.

Using this analysis code, it can be seen that introducing a 3x3 training grid into the precipitation downscaling code significantly reduced pixilation for individual rain storms, creating spatial statistics for the downscaled data that resembled observed characteristics much more closely than the original fields. Figure 13 compares the characteristics of the differences between each individual grid cell for observed (top), model old (second row), model new (third row) for GFDL CM2.1. Similar results (not shown) were obtained for other models.

In Figure 13, the shape of the distribution is noticeably broader for the non-3x3 (old) downscaling (middle row) as compared to observations (top row), indicating a much wider range in inter-grid values. In the new 3x3 downscaling, the standard deviation of the distribution is significantly reduced and resembles observed much more closely.

The mean value of Moran's I for the "old" GFDL precipitation is 0.93 for a 25km radius and 0.69 for a 150km radius. The mean value for the "new" GFDL simulated precipitation is 0.97 for a 25km radius and 0.81 for a 150km radius, where the higher value indicates a spatial field that is more closely auto-correlated.

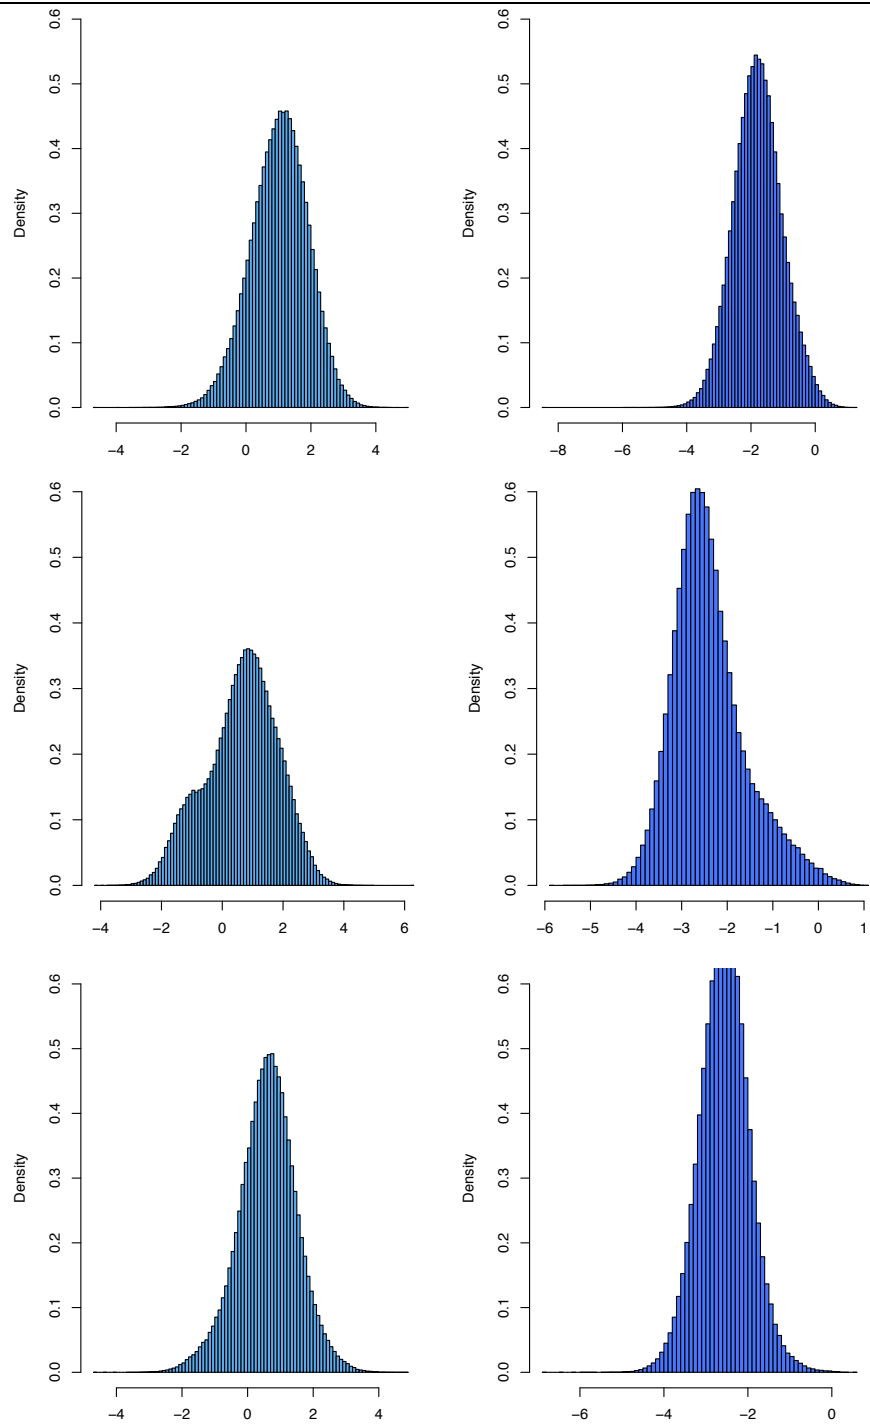
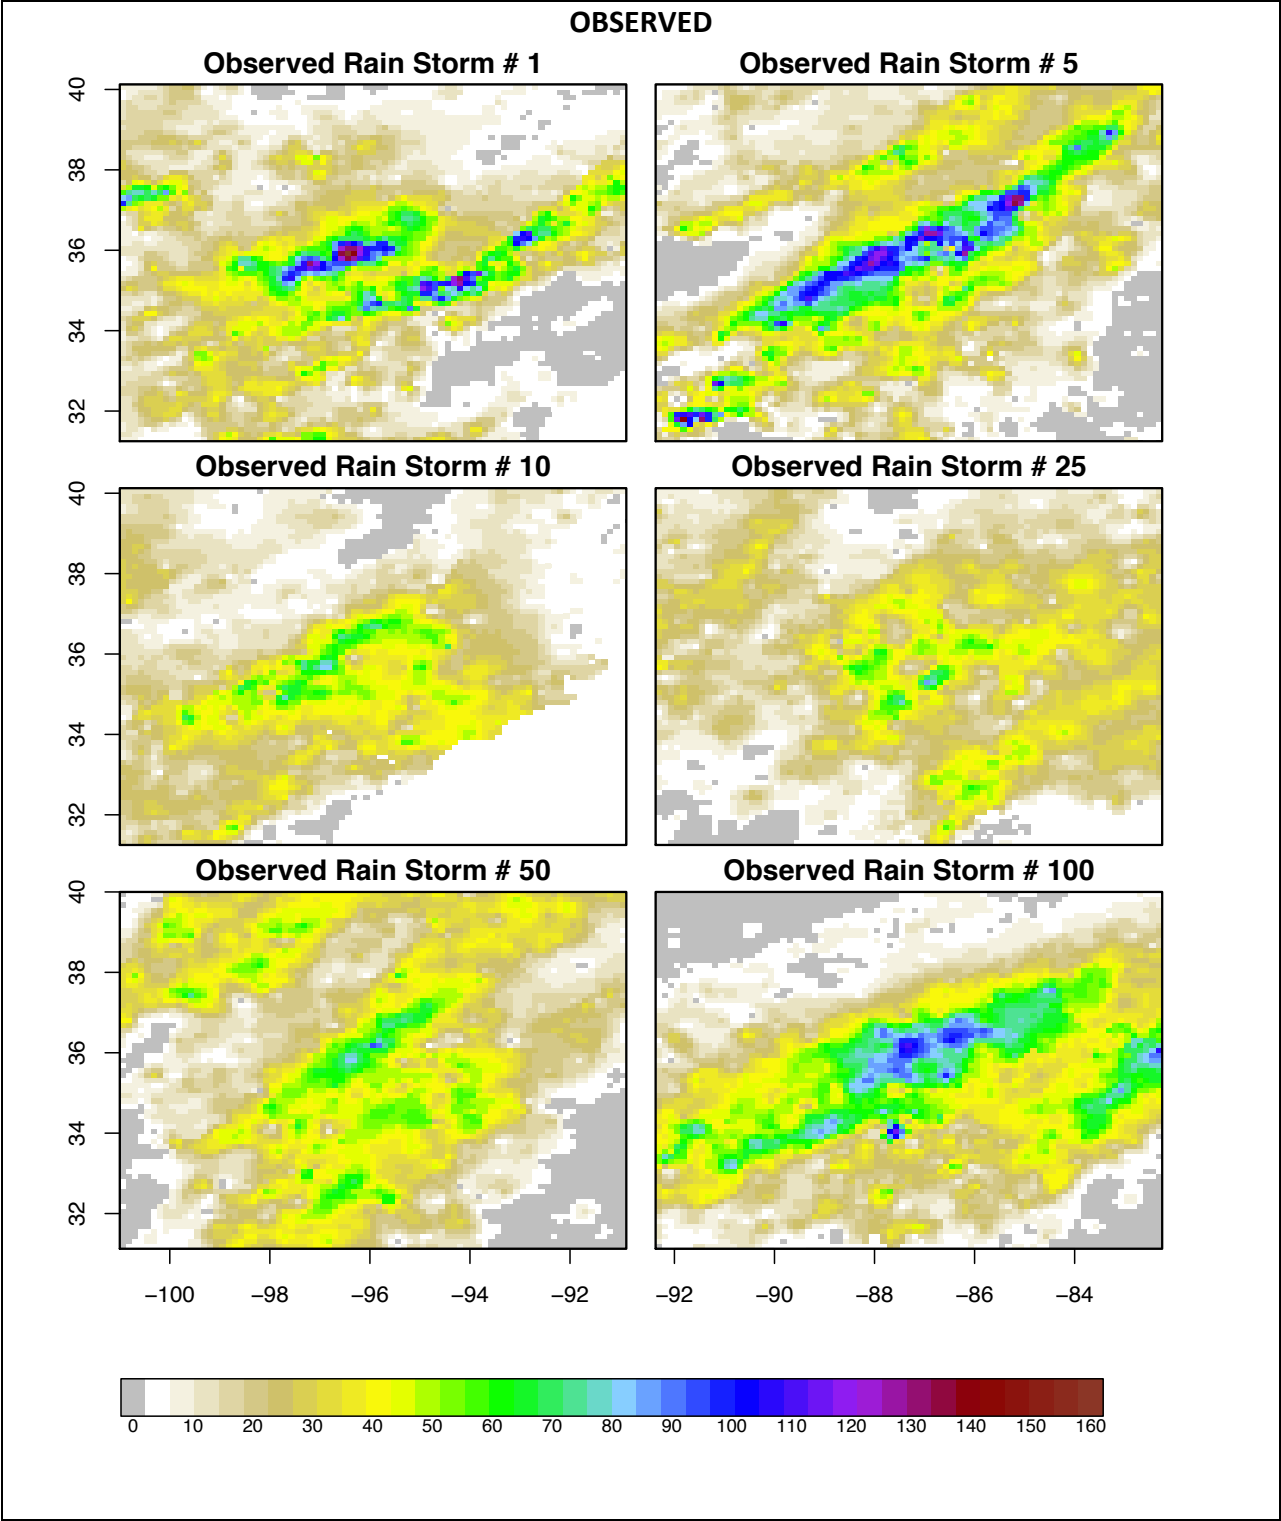
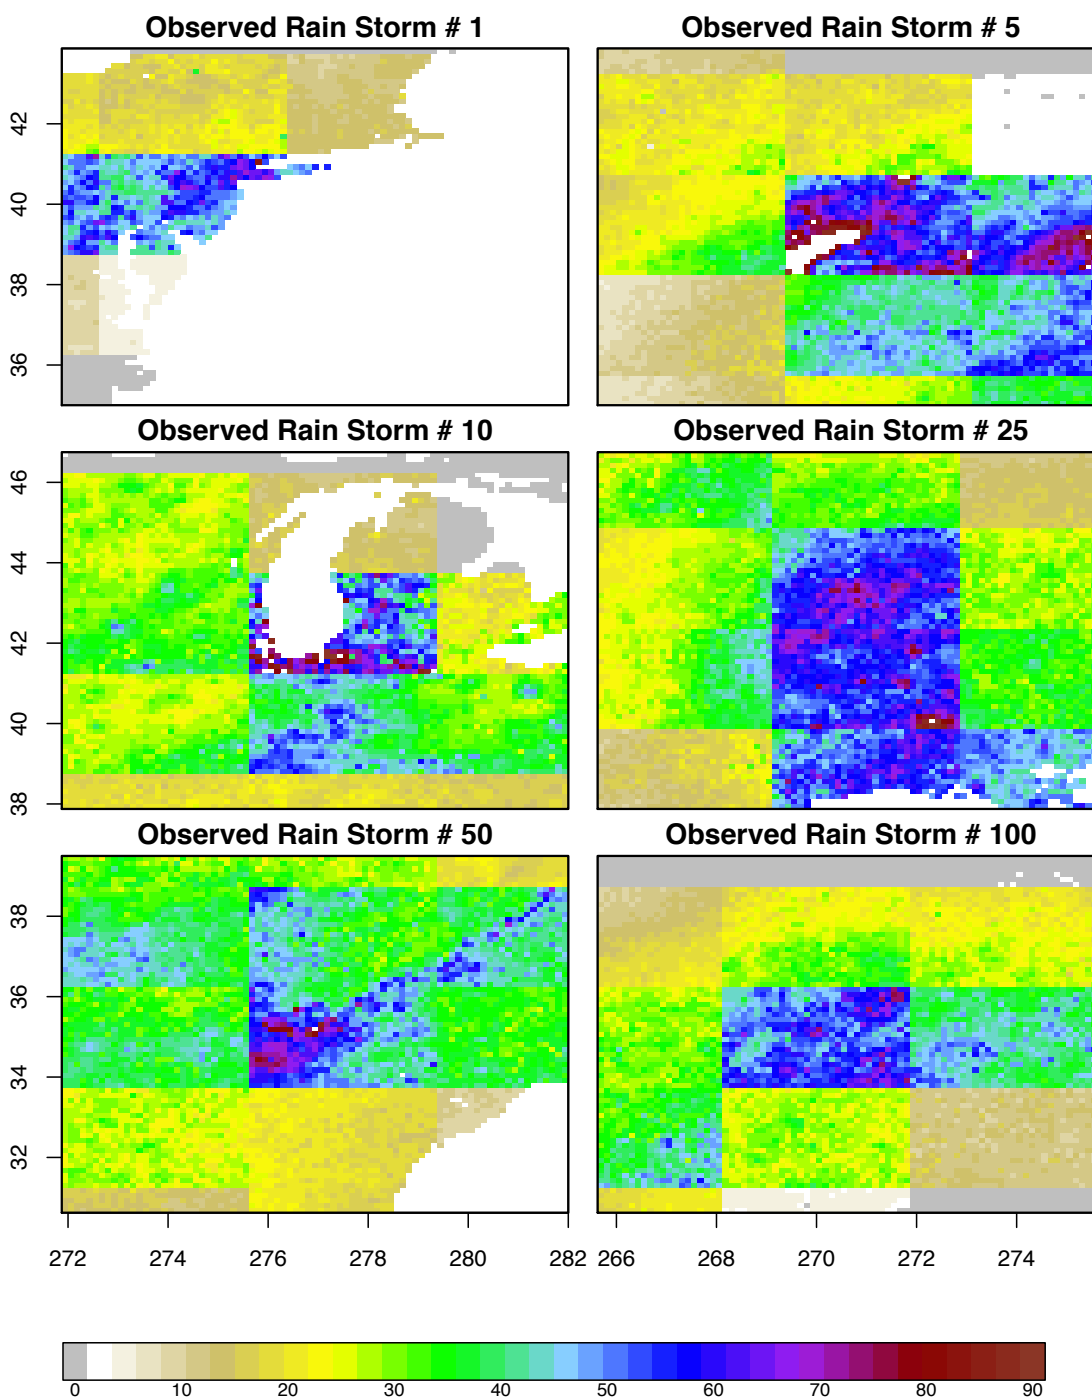


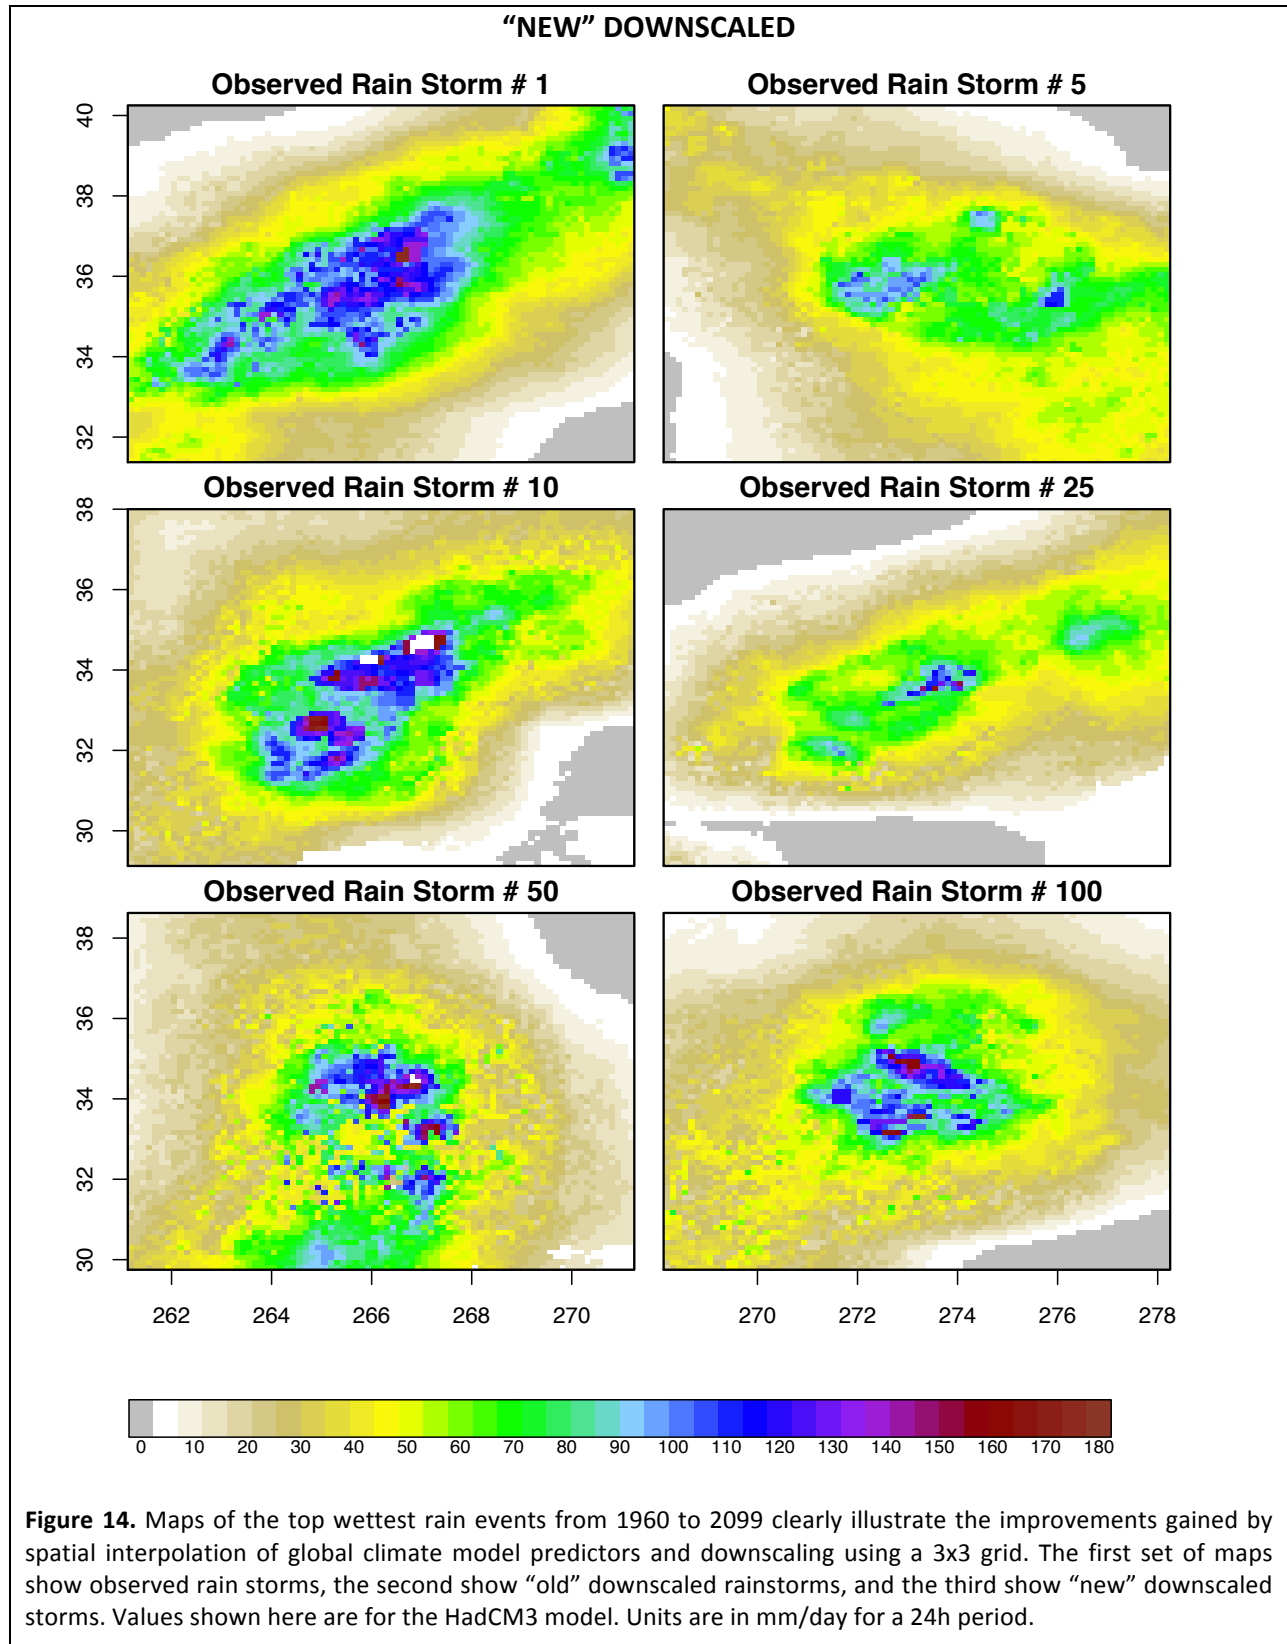
Figure 13. A histogram of the differences in 24h cumulative precipitation between four (up, down, left, right) neighboring grid cells (left) and eight (a 3x3 grid) neighboring grid cells (right) for the top 100 wettest precipitation events for the period 1960 to 1999, for observations (top), “old” downscaling where each grid cell was processed independently (middle) and “new” downscaling where grid cells were downscaled within a 3x3 matrix (bottom). Values shown here are for the GFDL CM2.1 model.

Figure 14 compares images of 24h cumulative precipitation resulting from six storms (the wettest, 5th, 10th, 25th, 50th and 100th) in the observed, old downscaled, and new downscaled dataset. This figure shows results from the HadCM3 model; similar results (not shown) were obtained for other models. “Old” downscaled precipitation showed a clear shadow of the original global climate model grid and was highly pixilated. “New” downscaled precipitation now exhibits a much more homogeneous structure that more closely resembles observed.



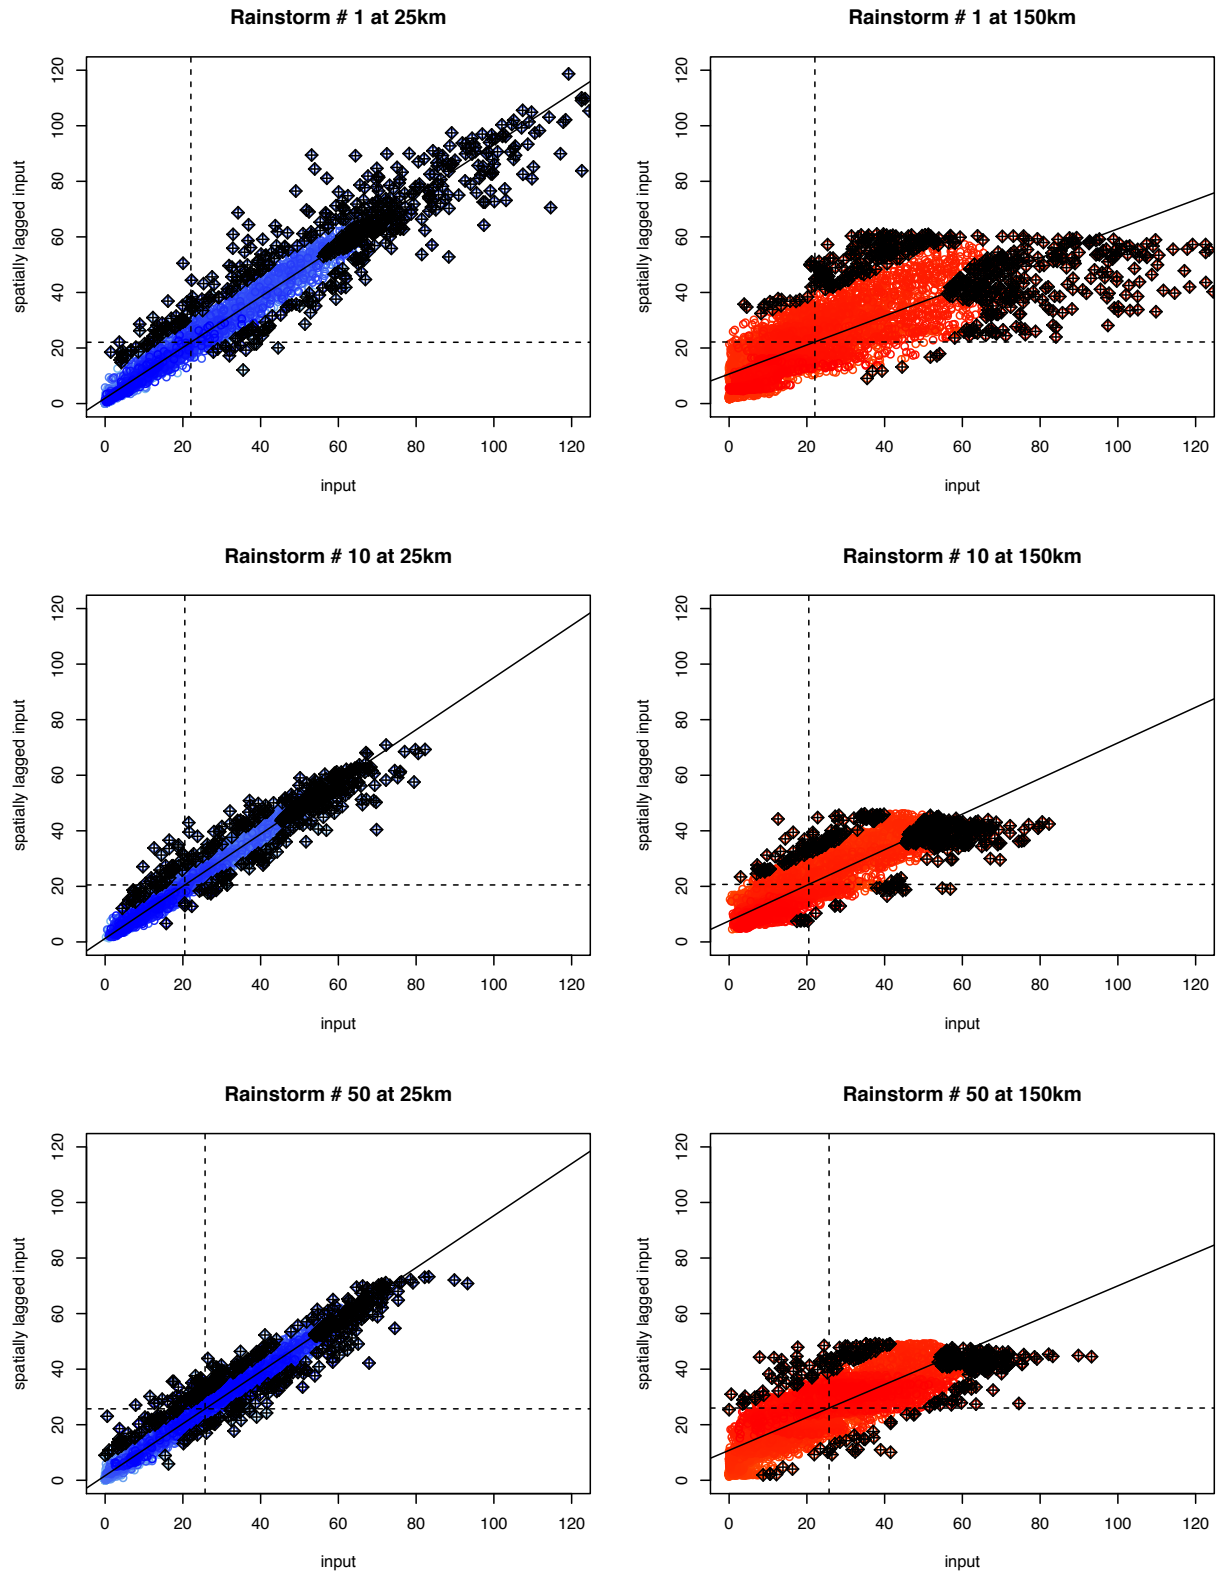
"OLD" DOWNSCALED






Finally, the value of Moran's I for individual grid cells is plotted and compared to observations in Figure 15 for spatial radii of 25km (left) and 150km (right). Results from the GFDL CM2.1 model are shown. Again, similar results were obtained for other models (not shown). While the spread in individual grid cell values is still higher than observed for the "new" downscaling as compared to the old, introducing the 3x3 downscaling clearly provides a significant improvement in the spatial autocorrelation of the precipitation fields for individual rain events.

OBSERVED




"OLD" DOWNSCALED


Rainstorm # 1 at 25km




Rainstorm # 1 at 150km




Rainstorm # 10 at 25km




Rainstorm # 10 at 150km

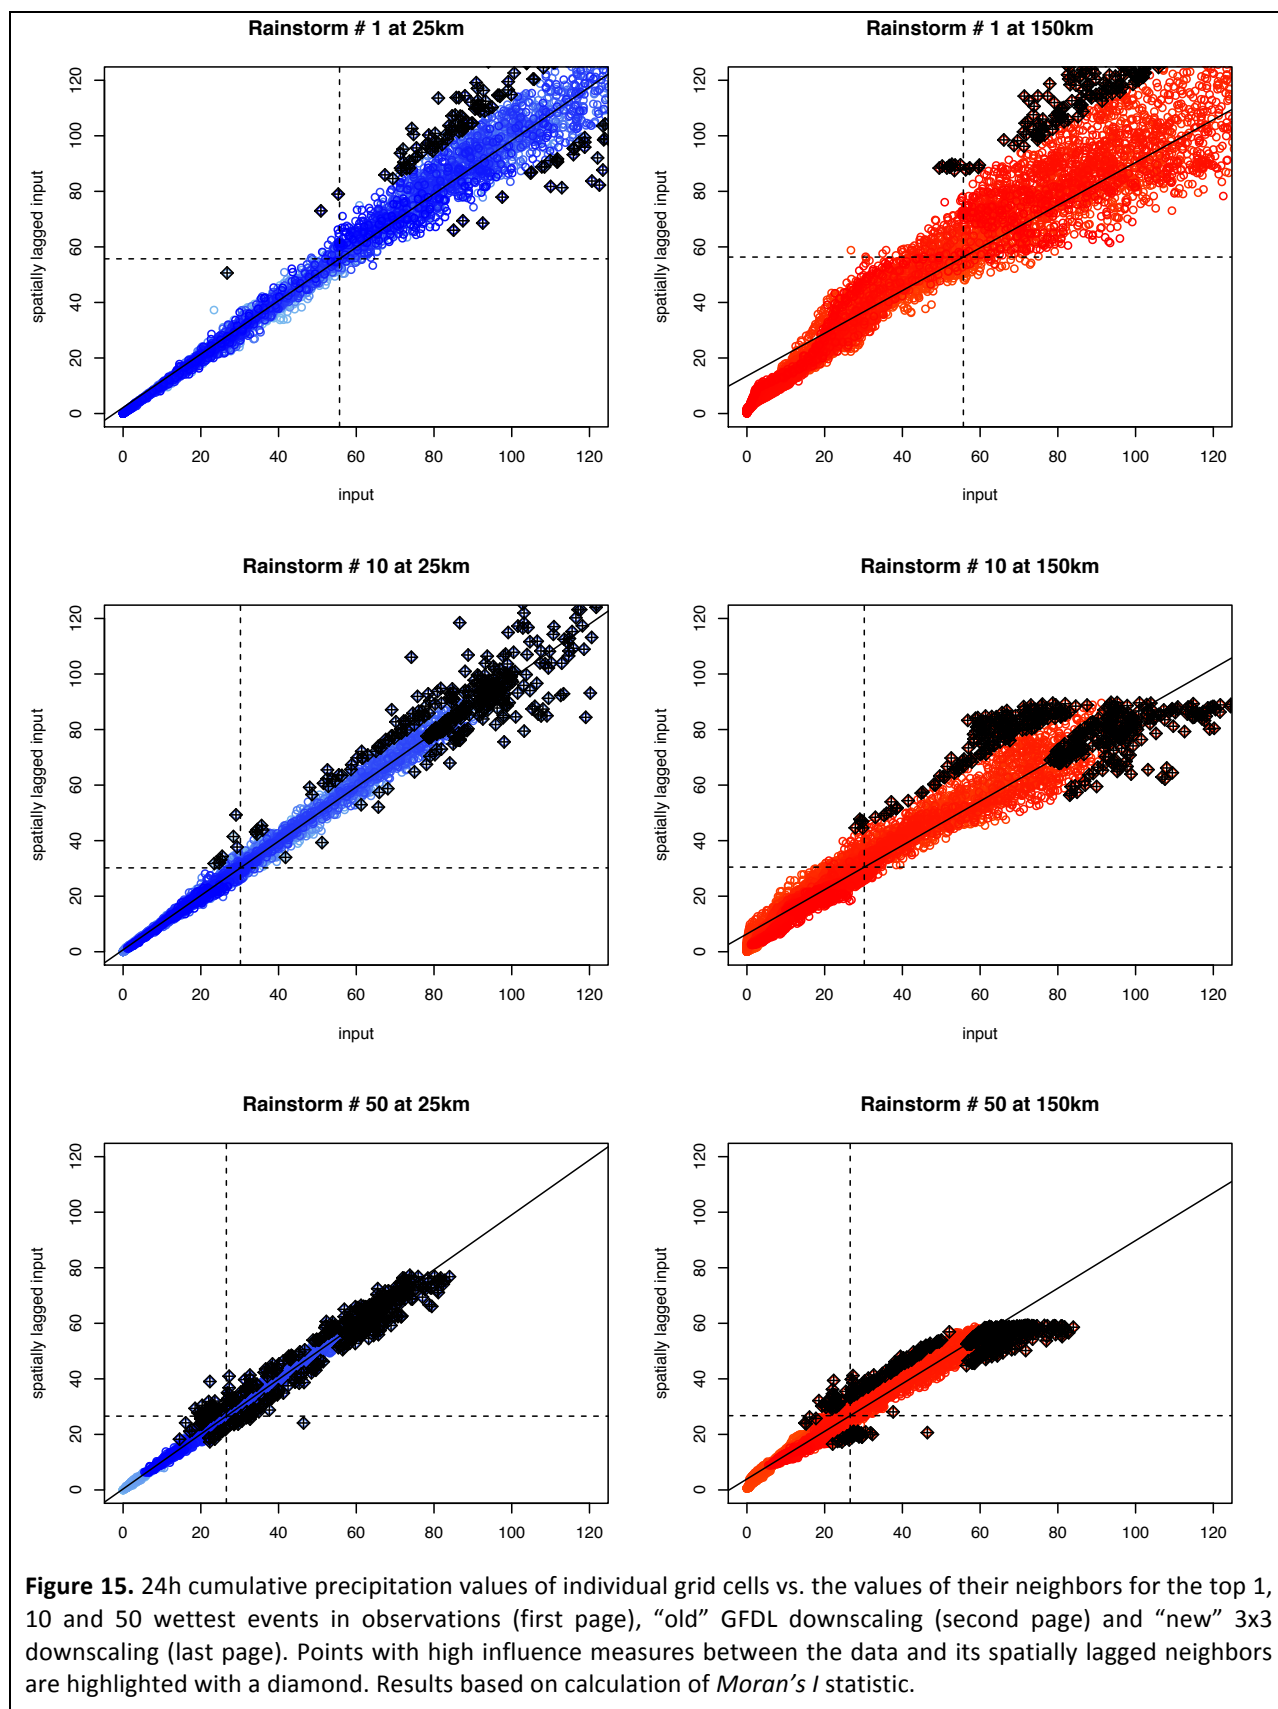


Rainstorm # 50 at 25km



Rainstorm # 50 at 150km





Improvements to Computation

With professional optimization and parallel computing resources from the TTU High Performance Computing Center, clock time to compute the full set of downscaling was reduced from more than one century to 1 year.

Optimization

At the beginning of the project, the downscaling code required approximately 20 minutes to complete calculations for one location. Multiplying this time by approximately 100,000 grid cells plus 10,000 individual station locations, 50 model/scenario combinations, and 3 variables yielded an estimated completion time of over 600 human years, assuming continuous run time. Although the original funding supported purchase of two additional servers, this only reduced the estimated run time to 200 years.

The first step in reducing the computational demand of this project was to identify an expert in optimization of R code. Dr. Matthew Pratola, formerly of Los Alamos National Laboratory, completed his dissertation on this topic. With additional research funds external to this project, he was engaged to optimize the downscaling code and was able to successfully reduce the run time for an individual location from an average of 20 minutes to 30 seconds. This reduced the scope of the project to 15 years on an individual server and 5 years using the additional two servers purchased for this project.

Parallelization

The next step in improving the computational efficiency of the downscaling was to parallelize the code to enable application using the Texas Tech University High Performance Computing Center (TTU-HPCC) resources as well as TTU shares in the Texas Advanced Computing Center hosted at the University of Texas.

The code was parallelized by dividing the observations and global climate model predictors up into individual longitude slices and providing one slice to each processor. Statistical downscaling code is nearly infinitely parallelizable, as it can be sliced into smaller and smaller pieces as needed. Too much parallelization, however, can be counter-productive as the number of files multiplies the time required to post-process the resulting “minifiles” into a region-wide grid using a series of scripts invoking netCDF NCO and CDO utilities. For that reason, a balance was struck between run times and file sizes. Computation time for 222 grid cells averaged between 2 to 4 hours for temperature (longer than on an individual server as I/O times were greatly reduced although computational time remained similar), slightly longer for precipitation with one predictor, and more than double for precipitation with multiple predictors.

Taken together, this optimization and parallelization of the code enabled all downscaling to be completed in approximately one year of computer clock time. Much of the original downscaling had to be re-calculated with EOF-filtered temperature, spatially interpolated predictors, and the new 3x3 precipitation code, which doubled the duration of the calculations and total CPU usage to more than 2 million over two years.

This optimized and parallelized computing framework will be used for future downscaling of CMIP5 simulations.

Error and Bias Analysis

A comprehensive error and bias analysis code evaluates 22 types of errors and 74 types of biases in the high-resolution climate projections relative to observations. This is a generalizable code that can be used to quality control any downscaled dataset or compare any observational datasets.

Error Tests

In conducting any analysis involving tens of thousands of files and multiple terabytes of data, likelihood of error is high. For that reason, we developed a rigorous framework to evaluate the downscaled projections and the gridded observations for 17 types of errors in temperature fields and 5 types of errors in precipitation. Errors include values exceeding the highest or lowest observed historical values for the continent and for that individual grid cell; missing values; repeated values; days where minimum temperature is higher than maximum temperature (as we knew this error already occurred in station data); and lack of expected seasonal differences or trends in the data over time.

The first iteration of these error tests identified a number of “minifiles” which had saved incorrectly, as well as peculiarities to various model inputs that resulted in erroneous missing values. These errors were corrected and the resulting final datasets incorporate those corrections. Each specific type of error is listed in the tables below and maps of the errors for each global model and gridded observed values are provided in **Appendix D**.

TEMPERATURE ERRORS (17 TESTS)

Test	Model	Obs
REALITY CHECK		
1. Daily maximum temperature value exceeds the continental US highest observed daily temperature	Pass/Fail for the historical period (simulated 1960-2009; observed 1960-1999)	
2. Daily maximum temperature value is lower than the continental US lowest observed daily temperature	“	
3. Daily minimum temperature value exceeds the continental US highest observed daily temperature	“	
4. Daily minimum temperature value is lower than the continental US lowest observed daily temperature	“	
EXTRAPOLATION CHECK		
5. Simulated daily maximum temperature exceeds highest daily maximum temperature observed at that location or grid cell (1960-1999)	Total number of days in simulated historical period (1960-2009)	N/A
6. Simulated daily maximum temperature is lower than the lowest daily maximum temperature observed at that location or grid cell (1960-1999)	“	N/A
7. Simulated daily minimum temperature exceeds highest daily minimum temperature observed at that location or grid cell (1960-1999)	“	N/A
8. Simulated daily minimum temperature is lower than the lowest daily minimum temperature observed at that location or grid cell (1960-1999)	“	N/A
CONSISTENCY CHECK		
9. Daily minimum temperature exceeds daily maximum temperature	Percentage of days per year (simulated 1960-2009; observed 1960-1999)	
MISSING VALUES CHECK		
10. Missing (NA) daily maximum temperature values	Total number of days in entire period (simulated 1960-2009; observed 1960-1999)	
11. Missing (NA) daily minimum temperature values	“	
TREND TEST		
12. Does maximum daily temperature increase from 1960-1989 to 2020-2049, 2050-2079, and 2080-2099?	Number of times it does not increase, from 0 to 3	N/A
13. Does minimum daily temperature increase from 1960-1989 to 2020-2049, 2050-2079, and 2080-2099?	“	N/A
REPEATS CHECK		
14. How many times do daily maximum temperature values repeat more than 3 times in a row (to 0.01°C)?	Total number of times in historical period (simulated 1960-2009; observed 1960-1999)	
15. How many times do daily minimum temperature values repeat more than 3 times in a row (to 0.01°C)?	“	
SEASON TEST		
16. How many times is the average of Jan-Feb daily maximum temperature warmer than the average of Jul-Aug for the same year?	Total number of times in period (simulated 1960-2009; observed 1960-1999)	
17. How many times is the average of Jan-Feb daily minimum temperature warmer than the average of Jul-Aug for the same year?	“	

PRECIPITATION ERRORS (5 TESTS)

Test	Model	Obs
REALITY CHECK		
1. 24h cumulative precipitation exceeds the continental US highest observed value	Pass/Fail for the total period (simulated 1960-2099; observed 1960-1999)	
2. 24h cumulative precipitation values below zero	“	
EXTRAPOLATION CHECK		
3. Simulated 24h precipitation exceeds highest 24h precipitation observed at that location or grid cell	Total number of days in historical period (1960-1999)	N/A
MISSING VALUES CHECK		
4. Missing (NA) daily precipitation values	Total number of days in entire period (simulated 1960-2099; observed 1960-1999)	
REPEATS CHECK		
5. How many times do non-zero 24h cumulative precipitation values repeat more than 3 times in a row (to 0.1mm)?	Total number of times in historical period (simulated 1960-2099; observed 1960-1999)	

Bias Tests

Biases in downscaled projections are typically evaluated using a cross-validation approach, where the statistical model is trained on one set of historical data (e.g. odd years) and then evaluated relative to an independent set of historical data (e.g. even years). Lacking future “observations”, this type of cross validation is primarily used to assess the stationarity of the model: i.e. to what extent is the model able to reproduce climate conditions not used to train the model? These types of biases also reflect the adequacy of the initial training data: i.e., did the sample of local climate used in training adequately capture all the important relationships between large-scale circulation at the spatial scale of the global climate model and high-resolution temperature and precipitation on the ground? Is the statistical model able to reproduce an independent set of climate conditions for that location?

As discussed previously, the ARRM method is currently being evaluated in a separate USGS-funded project that assesses the stationarity of this statistical downscaling method relative to future projections by end-of-century under higher emissions. In this project, ARRM performance is not compared to an independent sub-set of historical data, but rather to a set of high-resolution 25km dynamical model simulations with significantly different characteristics than the historical period on which the statistical model was trained. As shown in Figures 4-6, the ARRM method is relatively stationary even by end-of-century, with the exception of systematic warm biases along the coast for high quantiles of both maximum and minimum temperature.

Hence, rather than reevaluate the stationarity of ARRM-based downscaling relative to historical data, instead we chose to evaluate the ability of the ARRM model to remove global climate model biases, as this is the primary function statistical downscaling models are intended to perform.

To accomplish this goal, the ARRM model was trained on the full set of historical data from 1960 to 1999 then used to downscale to that same period (not an independent dataset) in order to quantify the ability of the ARRM model to correct for global climate model biases. It is important to note that a cross-validation approach would not resolve this specific feature aspect of the downscaling model, and equally important to note that this approach does provide any information on the stationarity of the downscaling method, as would a cross-validation approach. These are two separate questions, and we address them in two completely different ways.

Although the ARRM downscaling approach is able to remove much of the biases in global climate models at the regional to local scale, this bias analysis shows how the choice of global model still influences the resulting projections. Biases were calculated for 48 temperature and 26 precipitation indicators, including quantiles from 0.1 to 99.9, seasonal temperature and precipitation, and thresholds and extremes. For some models, the ARRM downscaling is able to remove much of the bias relative to historical observations. For other models (particularly for BCC-BCM2), the statistical model is not able to remove all of the biases due to the global model simulations.

Each specific type of bias that was examined is listed in the tables below and maps of the biases corresponding to each global model as compared to historical gridded observed values for the same time period (1960-1999) are provided in **Appendix E**.

TEMPERATURE BIASES (48 TESTS, all comparing simulated values to observations for the period 1960-1999, model minus observed values)

Test	Model
QUANTILES	
1-9. Values of the 0.1, 1, 10, 25, 50, 75, 90, 99, and 99.9 th quantiles of the distribution of daily maximum temperature	Absolute value, degrees C
10-18. Same, for daily minimum temperature	"
THRESHOLDS	
19-22. Number of days with daily maximum temperature above 32, 65, 90, and 100°F	Percentage of days, calculated as (mod-obs)/obs
23-26. Number of days with daily minimum temperature below 20, 32°F and above 65, 85°F	"
EXTREMES	
27. Average temperature of hottest 5-day period of the year (from daily maximum temperature)	Absolute value, degrees C
28. Average temperature of coldest 5-night period of the year (from daily minimum temperature)	"
29. Average temperature of hottest day of the year (from daily maximum temperature)	"
30. Average temperature of coldest night of the year (from daily minimum temperature)	"
31. Temperature of the hottest day in 30 years (from daily maximum temperature)	"
32. Temperature of coldest day in 30 years (from daily maximum temperature)	"
33. Temperature of warmest night in 30 years (from daily minimum temperature)	"
34. Temperature of coldest night in 30 years (from daily minimum temperature)	"
RANGE	
35. Difference between hottest day and coldest night in 30 years	Absolute value, degrees C
36. Difference between mean maximum and minimum temperature over 30 years	"
SEASONS	
37-40. Seasonal mean daily maximum temperature (winter or DJF; spring or MAM; summer or JJA; fall or SON)	Absolute value, degrees C
41-44. Same, for mean nighttime minimum temperature	"
IMPACT-RELEVANT INDICATORS	
45. Annual cumulative degree-days (threshold of 65°F)	Percentage of degree-days, calculated as (mod-obs)/obs
46. Onset of spring (defined as last day with minimum nighttime temperature below freezing)	Absolute value, days
47. Length of growing season (from the last spring frost to the first fall frost)	"
48. Onset of summer (defined somewhat arbitrarily as the first day of the year over 90°F)	"

PRECIPITATION BIASES (26 TESTS, all comparing simulated values to observations for the period 1960-1999, model minus observed values)

Test	Model
QUANTILES	
1-9. Values of the 0.1, 1, 10, 25, 50, 75, 90, 99, and 99.9 th quantiles of 24h cumulative precipitation (wet days only)	Absolute value, mm
THRESHOLDS	
10-11. Number of wet days with 24h cumulative precipitation above 0 (including trace precipitation) and above 0.01 (excluding trace precipitation)	Percentage of days, calculated as (mod-obs)/obs
12-15. Number of days with 24h cumulative precipitation above 0.5, 1, 2, and 3 inches	“
EXTREMES	
16. Average 24h cumulative precipitation on the wettest day of the year	Percentage of mm, calculated as (mod-obs)/obs
17. Average 24h cumulative precipitation on the 5-day wettest period of the year	“
18. 24h cumulative precipitation on the wettest day in 30 years	“
SEASONS	
19-22. Seasonal cumulative precipitation (winter or DJF; spring or MAM; summer or JJA; fall or SON)	Percentage of mm, calculated as (mod-obs)/obs
IMPACT-RELEVANT HYBRID INDICATORS	
23. Number of snow days per year	Absolute value, number of days
24. Ratio of precipitation falling as rain to that falling as snow	Difference in the ratio (mod-obs), mm/mm
25. Number of hot, dry days per year (defined as precipitation < 0.01” and daily maximum temperature > 90°F)	Absolute value, number of days
26. Number of cool, wet days per year (defined as precipitation > 0.01” and daily maximum temperature < 65°F)	Absolute value, number of days

Educating the User and the Downscaler

Initial application of downscaled projections revealed issues with both downscaler understanding of user needs and user application of the high-resolution projections. These have been addressed by improving downscaling methodology and documentation, including preparing a brief “user’s guide” to accompany the data online.

A large proportion of downscaling is conducted by practitioners who have a detailed understanding of the specific uses to which their high resolution projections will be put. This dataset is unique in that it is intended to be broadly used across a range of applications, many of which may be unfamiliar to the research team developing the projections.

Helpful input was provided by initial users of the dataset, who identified some issues with the data that had not previously been considered in the downscaling community. For example, a test for pixilation of daily rainfall is not part of any downscaling evaluation framework, but is obviously relevant to use of data in impact analyses. This and other related issues were incorporated back into the downscaling process, and where needed, original projections were re-generated in order to address user concerns. Also, a set of indicator analyses were developed, some of which were relevant to evaluating the dataset but others of which were primarily intended to provide a useful summary of the data to users in various impact communities.

Users also provided helpful feedback on *perceived* errors in the dataset, leading to the suggested creation of a “users guide” to high-resolution climate projections to adjust user expectations regarding what is and is not considered to be an error or a problem in climate projections. For example, users may assume that conditions on a given simulated day, such as July 18, 1992, should match observed; or that the presence of an “NA” value in a downscaled dataset is an error. The brief user’s guide on the next page provides general information relevant to the application and use of any climate model simulations in impact assessments.

A DUMMIES GUIDE TO HIGH-RESOLUTION CLIMATE PROJECTIONS

There is no one perfect climate model – so don't waste your time trying to find it! The scientific literature is clear that, for most analyses, the best approach is to try and use as many climate models as possible.

If an analysis of the climate model's ability to simulate large-scale weather patterns or other climate features that a model is *supposed to* be able to simulate at the scale of 100s to 1000s of km has shown that a certain model or models are demonstrably incapable of simulating relevant large-scale climate features over your region of interest (as is the case for models in the Arctic, for example – see Overland et al. 2011), then a case can be made for removing poorly-performing models from the ensemble.

Do not attempt to select a “best” model by comparing biases in seasonal temperature or precipitation for your location of interest. This evaluation method offers no guarantee that the top-performing models in terms of historical biases will be able to simulate the impacts of global change on the same region.

If using multiple climate model simulations for your analysis, always average across climate models as the very last step in the analysis. Unless the relationships between climate change and the impacts being studied are entirely linear, averaging across climate models too early in the analysis will remove the variability from the climate projections, leading to incorrect results.

Do not average across multiple emission scenarios. Scenarios are not like physical systems, where averaging can improve the quality of the data. Scenarios are entirely separate, independent pictures of what the future may look like, given a set of assumptions regarding socio-economic and technological development. Results of any analysis should be averaged across climate models, but presented independently for each emission scenario used.

There is no one most likely emissions scenario: not the highest, not the lowest, and certainly not the middle scenario. The one thing we know for sure is that it is impossible to predict human behavior. In fact, by studying the impacts of climate change, we can hope that we are dynamically changing the likelihood of these scenarios by our work!

Which scenario(s) to use depends on which questions are most relevant to your analysis. Using a higher scenario will quantify the impacts of continued dependence on fossil fuels. Using a lower scenario will capture the impacts that will have to be adapted to, even if mitigation occurs. Comparing the results from a higher vs. lower future scenario can quantify the benefits of mitigation in terms of impacts avoided.

Downscaled projections have missing or NA values in them for various reasons. Sometimes the global climate model outputs got messed up or misplaced. Some models only have 360 days per year. Some models only saved certain years. NA values are not errors; they are just characteristics of the data that have to be worked around.

Climate projections are intended to match observations over climate time scales of decades, not days. Do not expect a downscaled climate simulation to match day-to-day observations at any given location. The averages should match over 20-30 years, but climate models are allowed to develop their own unique patterns of day to day climate variability.

It's a good idea to use multiple global climate model simulations and multiple future emission scenarios, because these capture a range of scientific and socio-economic uncertainty in future projections. It is not necessarily a good idea to use multiple statistical downscaling methods. Differences between downscaling methods are typically not due to scientific uncertainty, but rather due to limitations in the way the various models resolve the full distribution of the variable of interest.

Over the next decade or two, the most important source of uncertainty in future projections is natural variability. This is due to the inertia of the climate system in responding to human emissions, the inertia of the socio-economic system in initiating change, and the fact that many major sources of natural variability in climate, which affect both global and regional climate, operate over time scales shorter than 20 years.

Scientific uncertainty in the response of the climate system to human activities is the main source of uncertainty in temperature over the next few decades, and in precipitation through the end of the century.

Human or scenario uncertainty (what our emissions will be) is the largest source of uncertainty in temperature past mid-century.

CONCLUSIONS AND RECOMMENDATIONS

This project has produced a 6-terabyte database of high-resolution projections of daily maximum and minimum temperature and 24-hour cumulative precipitation covering the continental U.S. and Alaska. Future projections are based on simulations from 16 global climate models and four future emission scenarios, from higher to lower, and have been generated for both a regular grid and for individual station locations for the period 1960 to 2099.

Many challenges were encountered during the course of this project. They included scale-related issues in completing the magnitude of calculations required to generate the database, new demands being made of the downscaled projections that required re-thinking the methods used to generate and evaluate the database, and having to quality-control both model and observational data that originated from external sources which had represented the data as being of sufficient quality.

In order to address these and other issues that arose in the course of the project, the scope of the project was significantly expanded to include 11 new tasks not included in the original funding request, as well as to re-generate a large amount of the original data including all of the precipitation downscaling. New tasks covered a broad range of topics, from professional optimization of the downscaling code to construction of a detailed quality control code for observed station data.

Despite the challenges and setbacks encountered during the course of this project, and despite the fact that overall project workload was increased (in terms of new analysis, new code developed) by more than a factor of 3 compared to the initial scope of work, every project task has been completed successfully and the results of this project are not limited to the actual data produced. The modeling framework developed and refined in this project now exists to be built up and applied to new climate model simulations and new frameworks.

At this time, we continue to evaluate and improve downscaling methods. We are involved in a number of efforts to interact with and educate users of climate data as well as downscalers (see below). We are already adding to the existing dataset with new CMIP5 simulations. As this data is made available for public use, we also recommend that it incorporate user ability to define and calculate secondary indicators and extract data for individual locations and time periods. We hope this project will provide the basis for continued efforts to standardize and evaluate downscaling methods and to build state-of-the-art high-resolution climate modeling datasets for impact assessments.

OUTREACH

The data and methods used in this database have been incorporated into presentations made to a broad range of professional audiences, from non-profit organizations such as the National Wildlife Federation to federal agencies such as the Department of Transportation or Fish and Wildlife Service to academic collaborators in the Climate Science Center network.

PI Hayhoe has authored a guidebook and a series of educational videos for the US FWS on use of climate projections in impact assessments that references this work, to be released in 2013.

The projections generated by this project are expected to be broadly used throughout the Climate Science Center network in the South Central region and at the national scale. They will also be used in the upcoming 2014 U.S. National Climate Assessment and in the Environmental Protection Agency's work on climate indicators.

This research has resulted in one published article (Stoner et al. 2012), one submitted article (a review article for the National Climate Projections Platform, submitted to EOS), and four articles in preparation, describing: (1) the station-based quality control process and results of its application to the GHCN and MIDAS datasets; (2) analysis of the database projections; (3) analysis of future global mean temperature-based projections; (4) and a discussion of the sources of uncertainty in downscaling. The work has also been presented at professional meetings such as the Fall Meeting of the American Geophysical Union.

REFERENCES

- Covey, C., K. M. AchutaRao, U. Cubasch, P. Jones, S. J. Lambert, M. E. Mann, T. J. Phillips, and K. E. Taylor (2003) An overview of results from the Coupled Model Intercomparison Project, *Global and Planetary Change*, 37(1-2) 103-133.
- Dettinger, M. D., D. R. Cayan, M. K. Meyer, and A. E. Jeton (2004) Simulated Hydrologic Responses to Climate Variations and Change in the Merced, Carson, and American River Basins, Sierra Nevada, California, 1900–2099, *Climatic Change*, 62(1-3) 283-317.
- Hay, L. E. (2000) A comparison of delta change and downscaled GCM scenarios for three, *Journal of the American Water Resources Association*, 36(2) 387-397.
- Hayhoe, K., C. Wake, B. Anderson, X.-Z. Liang, E. Maurer, J. Zhu, J. Bradbury, A. DeGaetano, A. Hertel, and D. Wuebbles (2008) Regional Climate Change Projections for the Northeast U.S., *Mitigation and Adaptation Strategies for Global Change*, 13, 425-436.
- Hayhoe, K., et al. (2004) Emissions pathways, climate change, and impacts on California, *Proceedings of the National Academy of Sciences of the United States of America*, 101(34) 12422-12427.
- Hayhoe, K., K. Dixon, A. Stoner, J. Lanzante and A. Radhakrishnan. (2012) Is the past a guide to the future? Evaluating the assumption of climate stationarity in statistical downscaling. *Presented at the 2012 Fall Meeting of the American Geophysical Union*.
- Knutti, R. (2010) The end of model democracy? *Climatic Change*, 102, 395–404, doi:10.1007/s10584-010-9800-2
- Knutti, R., R. Furrer, C. Tebaldi, J. Cermak and G. Meehl. (2010) Challenges in combining projections from multiple climate models. *Journal of Climate* 23, 2739-2758
- Maurer, E. P., A. W. Wood, J. C. Adam, D. P. Lettenmaier, and B. Nijssen (2002) A long-term hydrologically based dataset of land surface fluxes and states for the conterminous United States, *Journal of Climate*, 15(22) 3237-3251.
- Meehl, G. A., C. Covey, T. Delworth, M. Latif, B. McAvaney, J. F. B. Mitchell, R. J. Stouffer, and K. E. Taylor. (2007) The WCRP CMIP3 multi-model dataset: A new era in climate change research, *Bulletin of the American Meteorological Society*, 88, 1383-1394.
- Nakicenovic, N. et al (2000). *Special Report on Emissions Scenarios: A Special Report of Working Group III of the Intergovernmental Panel on Climate Change*, Cambridge University Press, Cambridge, U.K., 599 pp.
- National Research Council Committee on Stabilization Targets for Atmospheric Greenhouse Gas Concentrations (2010) *Climate Stabilization Targets: Emissions, Concentrations, and Impacts over Decades to Millennia*. Cambridge University Press, 190 pp.
- Overland, J., M. Wang, N. Bond, J. Walsh, V. Kattsov and W. Chapman. (2011) Considerations in the selection of global climate models for regional climate projections: the Arctic as a case study. *Journal of Climate* 24, 1583-1597.
- Raisanen, J., L. Ruokolainen and J. Ylhäisi. (2010) Weighting of model results for improving best estimates of climate change. *Climate Dynamics* 35, 407-422.

- Randall, D.A., R.A. Wood, S. Bony, R. Colman, T. Fichefet, J. Fyfe, V. Kattsov, A. Pitman, J. Shukla, J. Srinivasan, R.J. Stouffer, A. Sumi and K.E. Taylor (2007) *Climate Models and Their Evaluation*. In: Climate Change 2007: The Physical Science Basis. Contribution of Working Group I to the Fourth Assessment Report of the Intergovernmental Panel on Climate Change [Solomon, S., D. Qin, M. Manning, Z. Chen, M. Marquis, K.B. Averyt, M. Tignor and H.L. Miller (eds.)]. Cambridge University Press, Cambridge, United Kingdom and New York, NY, USA
- Moss, R., Jae A. Edmonds, Kathy A. Hibbard, Martin R. Manning, Steven K. Rose, Detlef P. van Vuuren, Timothy R. Carter, Seita Emori, Mikiko Kainuma, Tom Kram, Gerald A. Meehl, John F. B. Mitchell, Nebojsa Nakicenovic, Keywan Riahi, Steven J. Smith, Ronald J. Stouffer, Allison M. Thomson, John P. Weyant & Thomas J. Wilbanks, (2010) The next generation of scenarios for climate change research and assessment. *Nature* 463, 747-756(11 February 2010)
- Stoner, A., K. Hayhoe, X. Yang and D. Wuebbles (2012) An asynchronous regional regression model for statistical downscaling of daily climate variables. *International Journal of Climatology*, DOI: 10.1002/joc.3603
- Taylor, Karl E., Ronald J. Stouffer, Gerald A. Meehl (2012) An Overview of CMIP5 and the Experiment Design. *Bull. Amer. Meteor. Soc.*, 93, 485–498.
- U.S. Global Change Research Program (USGCRP, 2009) Global Climate Change Impacts in the United States. Cambridge University Press. Available online at: <http://www.globalchange.gov/usimpacts/>.
- UK Meteorological Office. (2013) Met Office Integrated Data Archive System (MIDAS) Land and Marine Surface Stations Data (1853-current), [Internet]. NCAS British Atmospheric Data Centre, 2012. Available from http://badc.nerc.ac.uk/view/badc.nerc.ac.uk__ATOM__dataent_ukmo-midas
- U.S. Global Change Research Program (USGCRP, 2000) US National Assessment of the Potential Consequences of Climate Variability and Change. Cambridge University Press.
- Vrac, M., M. Stein, and K. Hayhoe (2007) Statistical downscaling of precipitation through nonhomogeneous stochastic weather typing, *Climate Research*, 34(3) 169-184.
- Wood, A. W., R. Leung, V. Sridhar and D. Lettenmaier (2004) Hydrologic implications of dynamical and statistical approaches to downscaling climate model outputs, *Climatic Change*, 62(1-3) 189-216.

APPENDIX A. Weather Stations

This Appendix lists the ID, latitude, and longitude of the weather stations for which sufficient data was available to downscale daily maximum and minimum temperature and precipitation. There are a total of 8,212 station records for maximum temperature, 8,176 for minimum temperature, and 10,272 for precipitation. Due to its length, this Appendix is provided as a separate Excel file.


APPENDIX B. Regional Temperature-Precipitation Scatter Plots

This Appendix provides scatter plots of projected changes in temperature and precipitation for 8 regions: Pacific Northwest, Southwest, Northern Great Plains, Southern Great Plains, Midwest, U.S. Northeast and Canadian Maritimes, Mid-Atlantic, and Southeast. Projected changes or anomalies are shown for seasonal and annual values as well as for the average number of hot/dry and cool/wet days per year for 2020-2039, 2050-2069 and 2080-2099 relative to 1960-1979. Each scenario is indicated by a different symbol, and two versions of each plot are provided, with and without global climate model labels on each point.


This plot style was originally created by J. Raisanen as part of the IPCC Third Assessment Report DDC archive. Thanks to Chris Anderson (Iowa State) for suggesting their inclusion here.

Scatter plots are also provided as Excel files in electronic format.


Northeast + Southern Ontario and Quebec




Northeast US/Canada Annual 2020-2039




Northeast US/Canada Annual 2050-2069




Northeast US/Canada Annual 2080-2099




Northeast US/Canada Winter 2020-2039




Northeast US/Canada Winter 2050-2069




Northeast US/Canada Winter 2080-2099




Northeast US/Canada Spring 2020-2039




Northeast US/Canada Spring 2020-2039




Northeast US/Canada Spring 2020-2039




Northeast US/Canada Summer 2020-2039




Northeast US/Canada Summer 2050-2069




Northeast US/Canada Summer 2080-2099




Northeast US/Canada Fall 2020-2039




Northeast US/Canada Fall 2050-2069




Northeast US/Canada Fall 2080-2099




Northeast US/Canada Annual 2020-2039




Northeast US/Canada Annual 2050-2069




Northeast US/Canada Annual 2080-2099




Mid-Atlantic




Mid-Atlantic States Annual 2020-2039




Mid-Atlantic States Annual 2050-2069




Mid-Atlantic States Annual 2080-2099




Mid-Atlantic States Winter 2020-2039




Mid-Atlantic States Winter 2050-2069




Mid-Atlantic States Winter 2080-2099




Mid-Atlantic States Spring 2020-2039




Mid-Atlantic States Spring 2020-2039




Mid-Atlantic States Spring 2020-2039




Mid-Atlantic States Summer 2020-2039




Mid-Atlantic States Summer 2050-2069




Mid-Atlantic States Summer 2080-2099




Mid-Atlantic States Fall 2020-2039




Mid-Atlantic States Fall 2050-2069




Mid-Atlantic States Fall 2080-2099




Mid-Atlantic States Annual 2020-2039



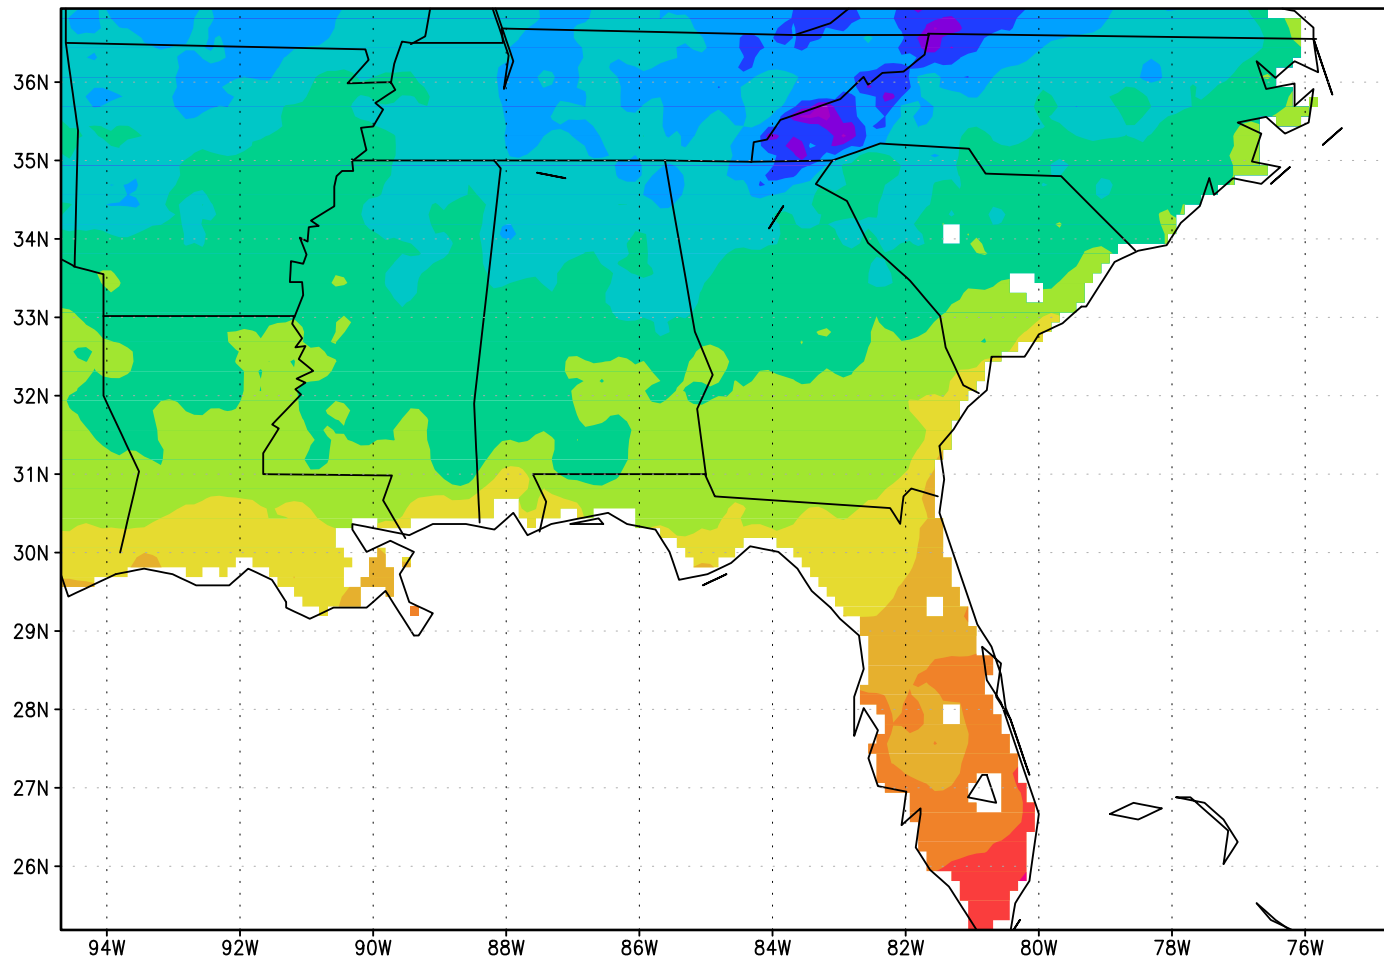
Mid-Atlantic States Annual 2050-2069




Mid-Atlantic States Annual 2080-2099




Southeast




Southeast Annual 2020-2039




Southeast Annual 2050-2069




Southeast Annual 2080-2099




Southeast Winter 2020-2039




Southeast Winter 2050-2069




Southeast Winter 2080-2099




Southeast Spring 2020-2039




Southeast Spring 2020-2039




Southeast Spring 2020-2039




Southeast Summer 2020-2039




Southeast Summer 2050-2069




Southeast Summer 2080-2099




Southeast Fall 2020-2039




Southeast Fall 2050-2069




Southeast Fall 2080-2099




Southeast Annual 2020-2039




Southeast Annual 2050-2069




Southeast Annual 2080-2099




Midwest




Midwest Annual 2020-2039




Midwest Annual 2050-2069




Midwest Annual 2080-2099




Midwest Winter 2020-2039




Midwest Winter 2050-2069




Midwest Winter 2080-2099




Midwest Spring 2020-2039




Midwest Spring 2020-2039




Midwest Spring 2020-2039




Midwest Summer 2020-2039




Midwest Summer 2050-2069




Midwest Summer 2080-2099




Midwest Fall 2020-2039




Midwest Fall 2050-2069




Midwest Fall 2080-2099




Midwest Annual 2020-2039




Midwest Annual 2050-2069




Midwest Annual 2080-2099




Northern Great Plains




Northern Plains Annual 2020-2039




Northern Plains Annual 2050-2069




Northern Plains Annual 2080-2099




Northern Plains Winter 2020-2039




Northern Plains Winter 2050-2069




Northern Plains Winter 2080-2099




Northern Plains Spring 2020-2039




Northern Plains Spring 2050-2069




Northern Plains Spring 2080-2099




Northern Plains Summer 2020-2039




Northern Plains Summer 2050-2069




Northern Plains Summer 2080-2099




Northern Plains Fall 2020-2039




Northern Plains Fall 2050-2069




Northern Plains Fall 2080-2099




Northern Plains Annual 2020-2039




Northern Plains Annual 2050-2069




Northern Plains Annual 2080-2099




Southern Great Plains




Southern Plains Annual 2020-2039




Southern Plains Annual 2050-2069




Southern Plains Annual 2080-2099




Southern Plains Winter 2020-2039




Southern Plains Winter 2050-2069




Southern Plains Winter 2080-2099




Southern Plains Spring 2020-2039




Southern Plains Spring 2050-2069




Southern Plains Spring 2080-2099




Southern Plains Summer 2020-2039




Southern Plains Summer 2050-2069




Southern Plains Summer 2080-2099




Southern Plains Fall 2020-2039




Southern Plains Fall 2050-2069




Southern Plains Fall 2080-2099




Southern Plains Annual 2020-2039




Southern Plains Annual 2050-2069




Southern Plains Annual 2080-2099




Pacific Northwest




Pacific Northwest Annual 2020-2039




Pacific Northwest Annual 2050-2069




Pacific Northwest Annual 2080-2099




Pacific Northwest Winter 2020-2039




Pacific Northwest Winter 2050-2069




Pacific Northwest Winter 2080-2099




Pacific Northwest Spring 2020-2039




Pacific Northwest Spring 2050-2069




Pacific Northwest Spring 2080-2099




Pacific Northwest Summer 2020-2039




Pacific Northwest Summer 2050-2069




Pacific Northwest Summer 2080-2099




Pacific Northwest Fall 2020-2039




Pacific Northwest Fall 2050-2069




Pacific Northwest Fall 2080-2099




Pacific Northwest Annual 2020-2039




Pacific Northwest Annual 2050-2069




Pacific Northwest Annual 2080-2099




Southwest




Southwest Annual 2020-2039




Southwest Annual 2050-2069




Southwest Annual 2080-2099




Southwest Winter 2020-2039




Southwest Winter 2050-2069




Southwest Winter 2080-2099




Southwest Spring 2020-2039




Southwest Spring 2020-2039




Southwest Spring 2020-2039




Southwest Summer 2020-2039




Southwest Summer 2050-2069




Southwest Summer 2080-2099




Southwest Fall 2020-2039




Southwest Fall 2050-2069




Southwest Fall 2080-2099




Southwest Annual 2020-2039



Southwest Annual 2050-2069



Southwest Annual 2080-2099

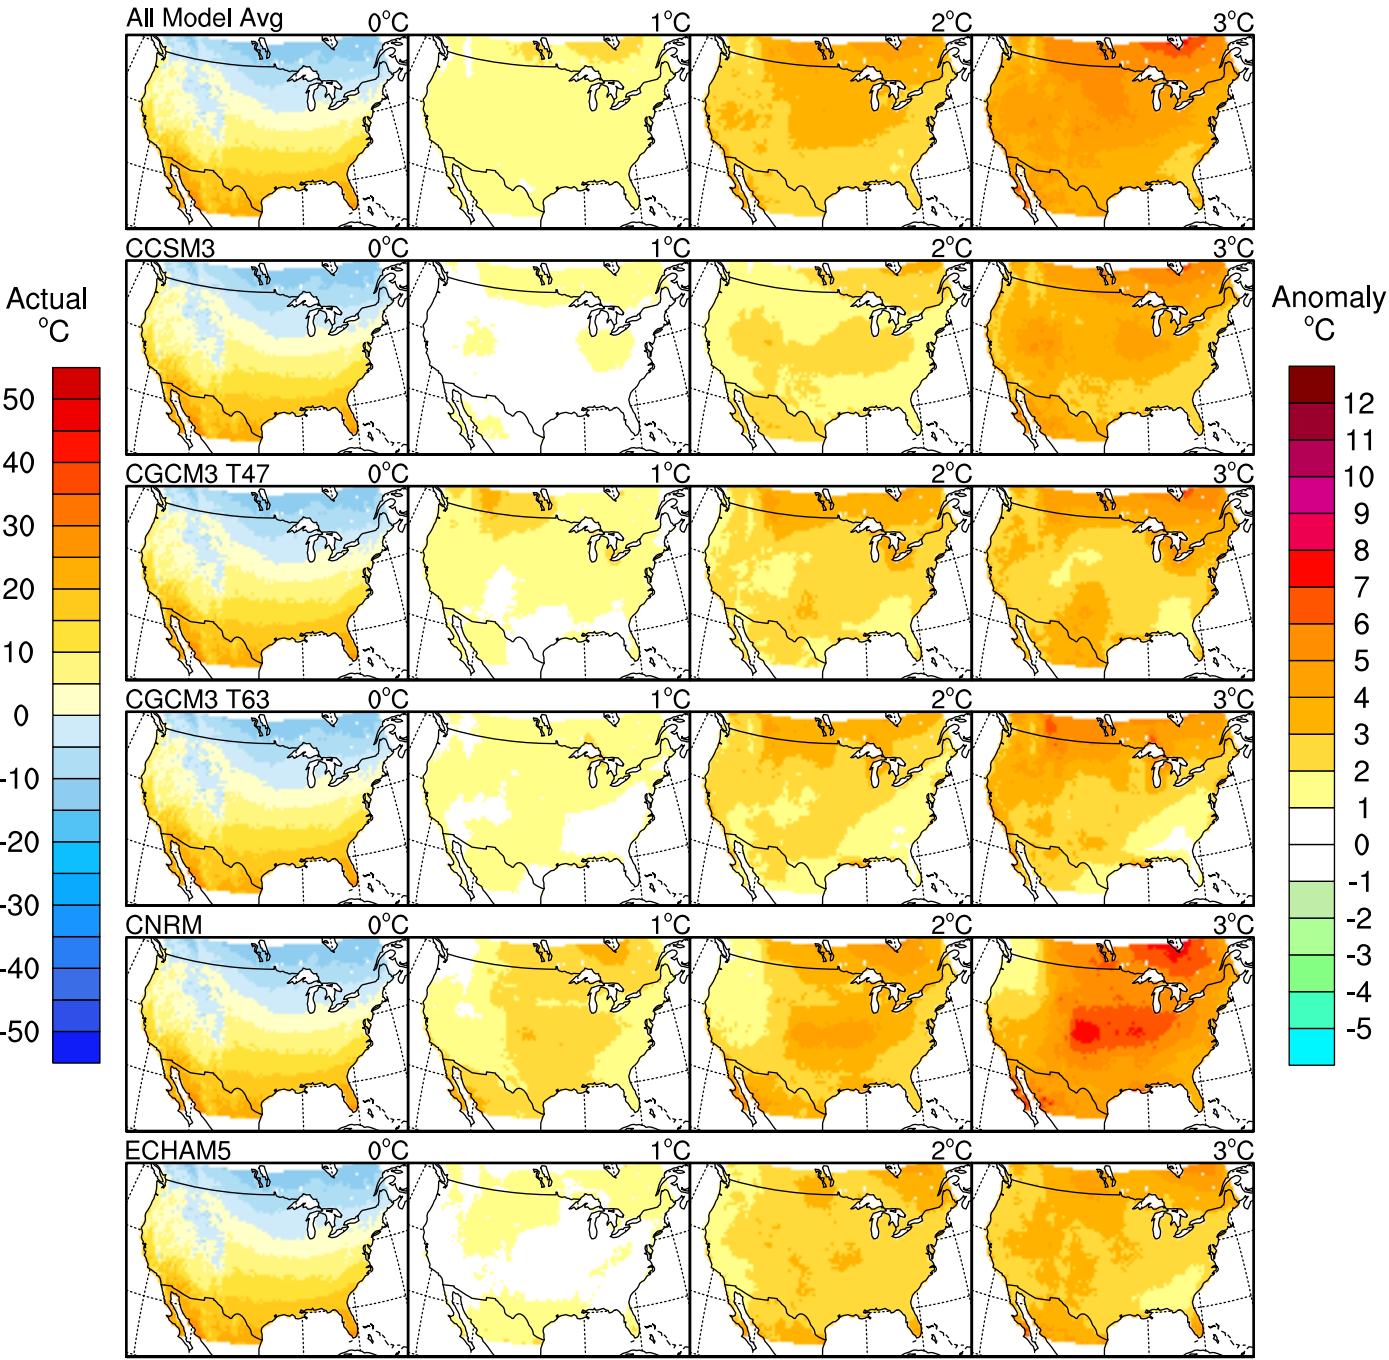


APPENDIX C. Climate Projections for Primary and Secondary Indicators

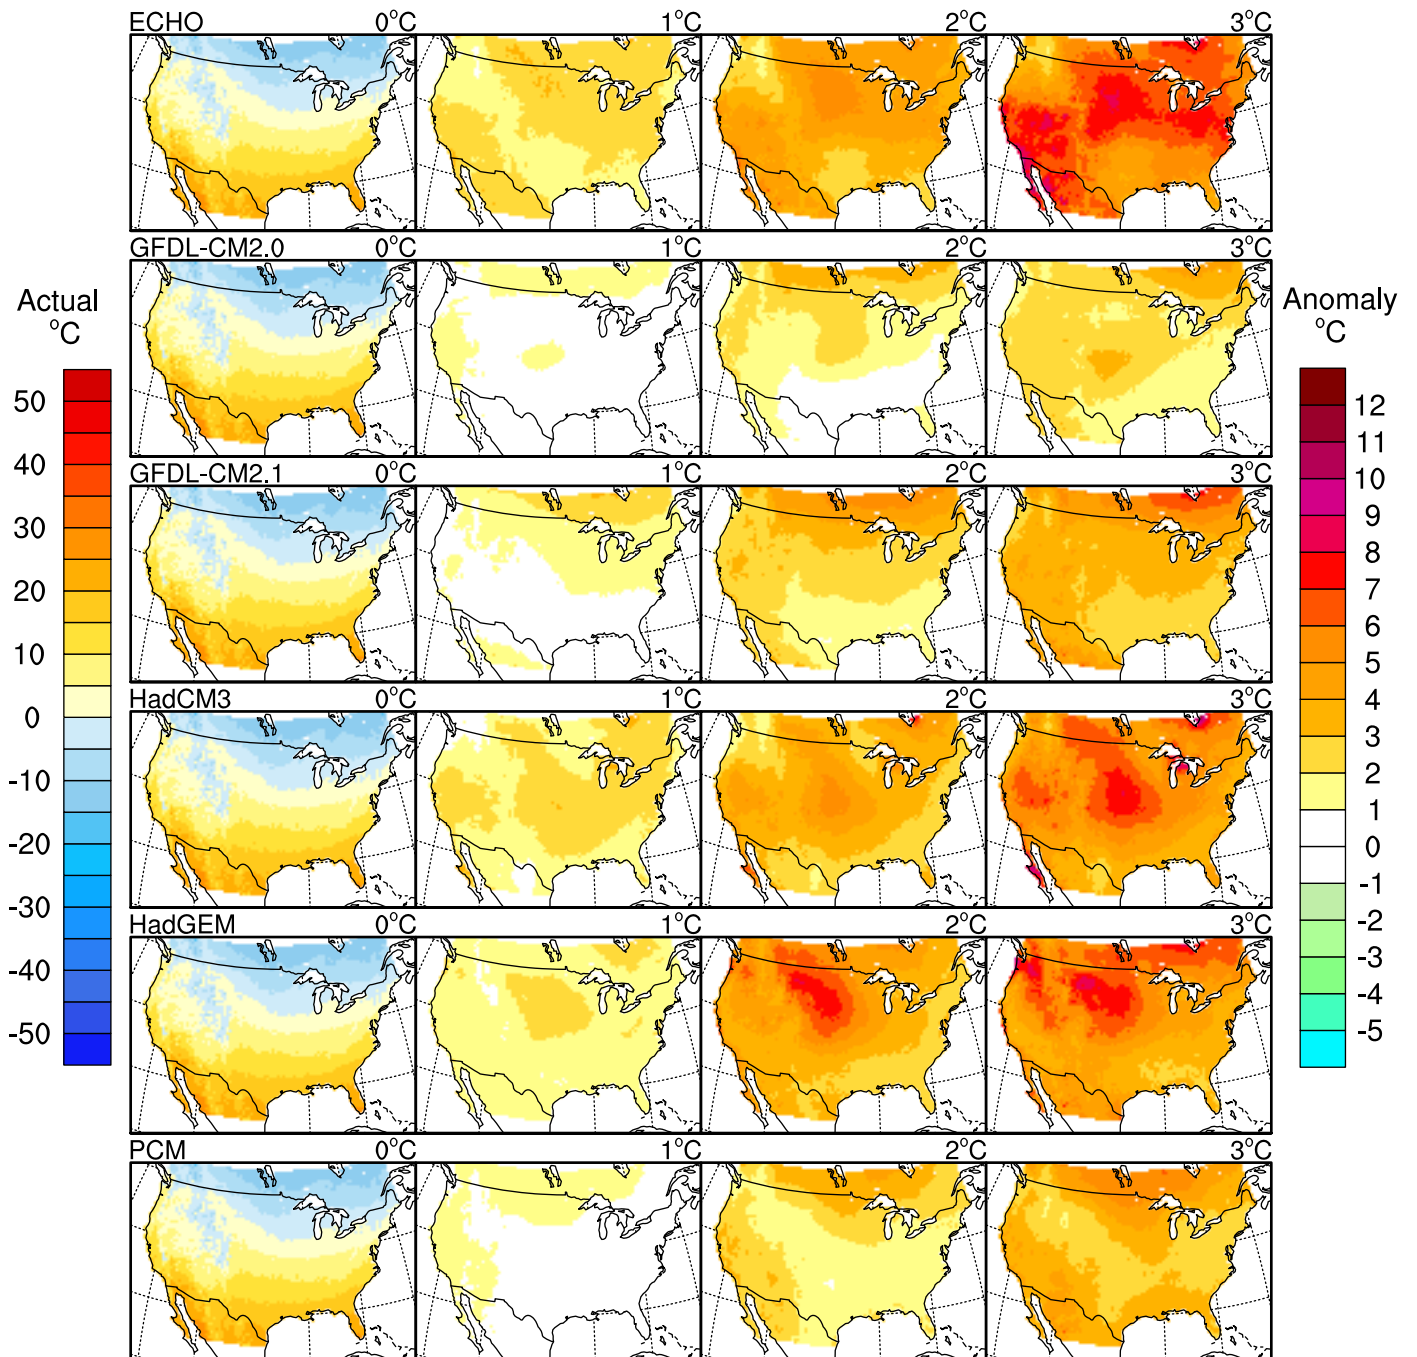
This Appendix provides maps of projected changes in a broad range of primary (temperature, precipitation) and secondary (thresholds, derived values) climate indicators as projected under +1, 2, and 3°C global mean temperature change relative to 1971-2000. The left-hand column of each set of maps provides the absolute value for 1971-2000, while the right-hand plots show the anomaly or difference relative to that period. The first row shows the all-model average, while subsequent rows show projected values for individual global climate models.

Original postscript files are also provided in electronic format.

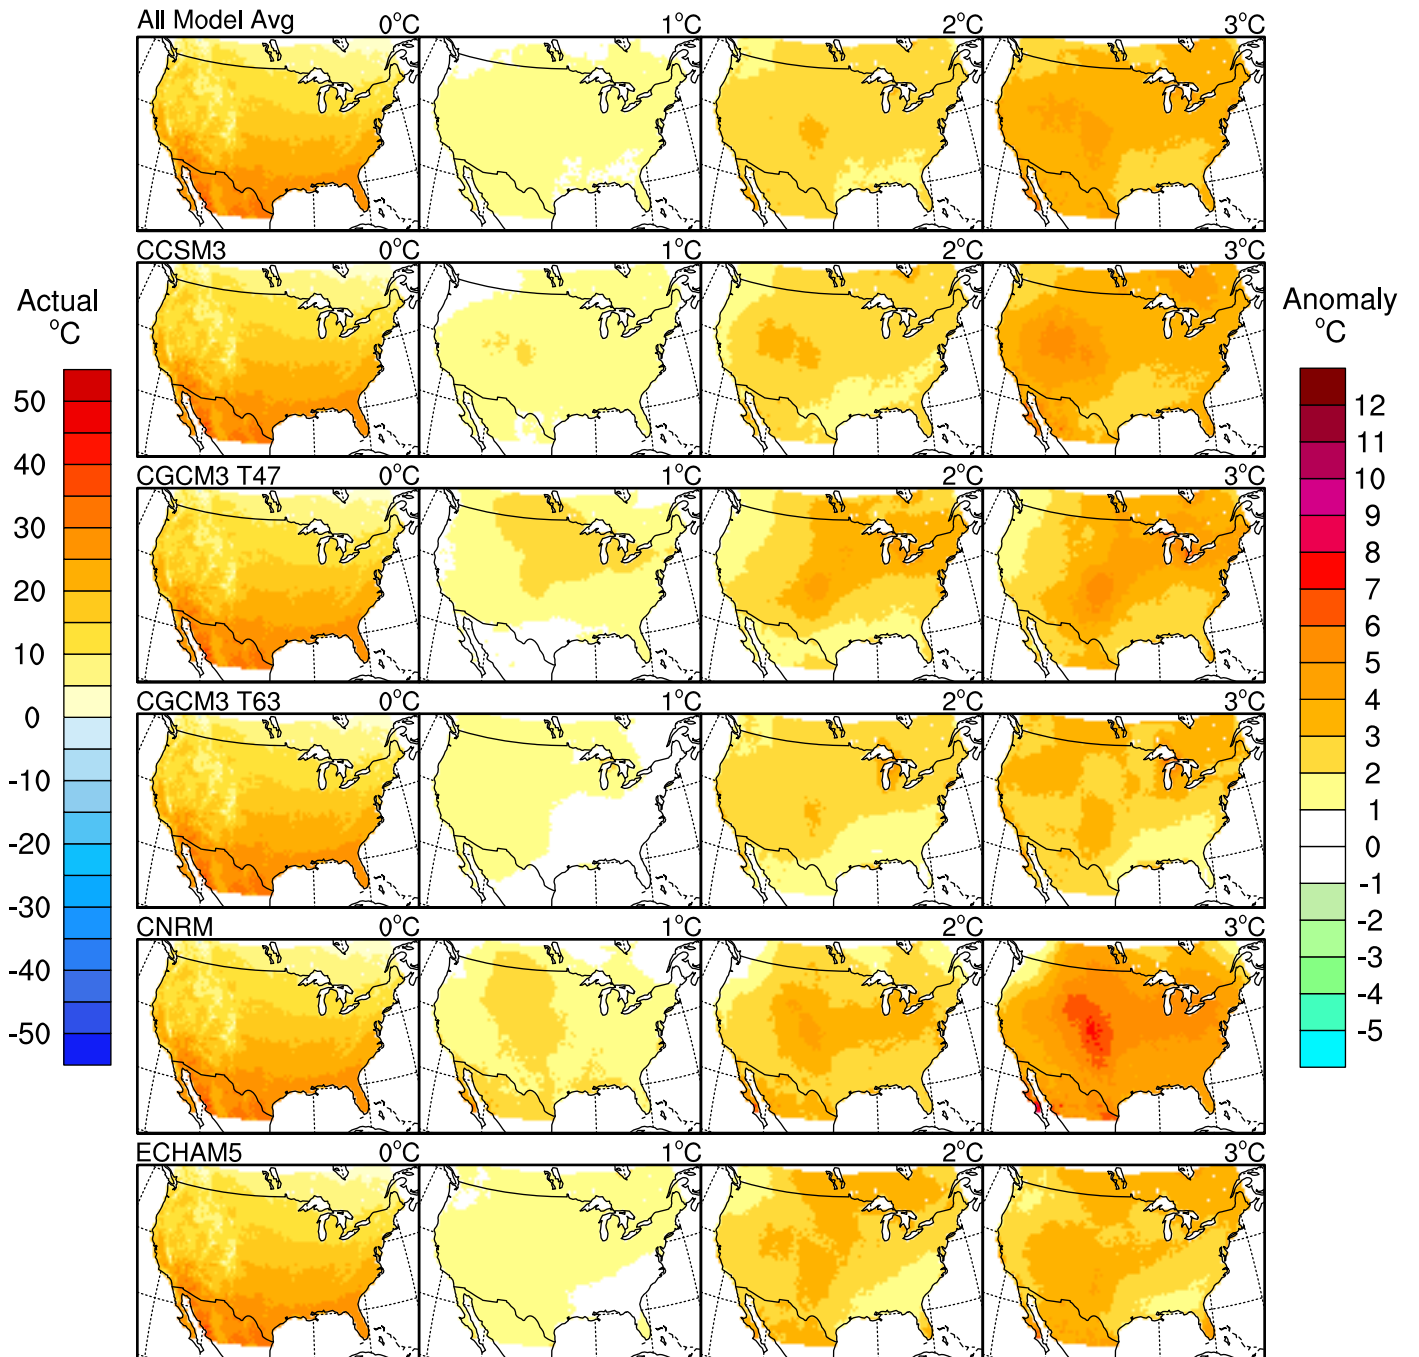
Avg Winter (DJF) Tmax




Avg Winter (DJF) Tmax




Avg Spring (MAM) Tmax



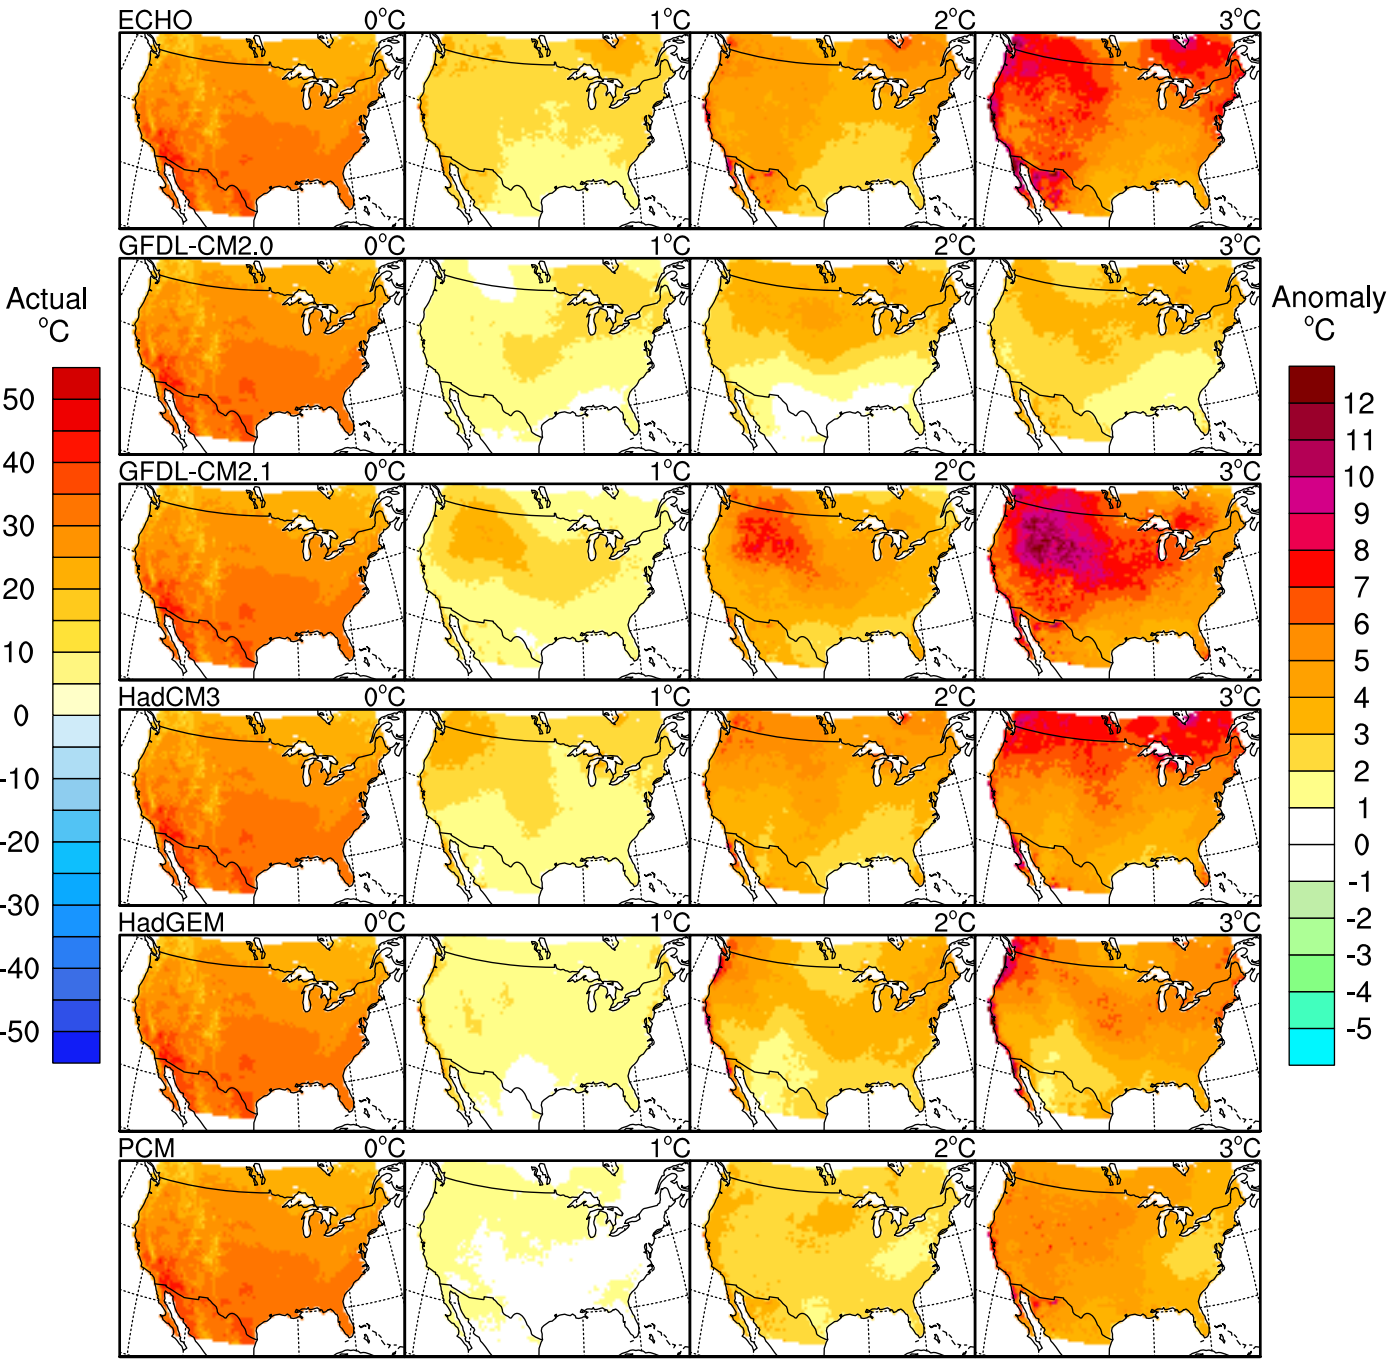
Avg Spring (MAM) Tmax



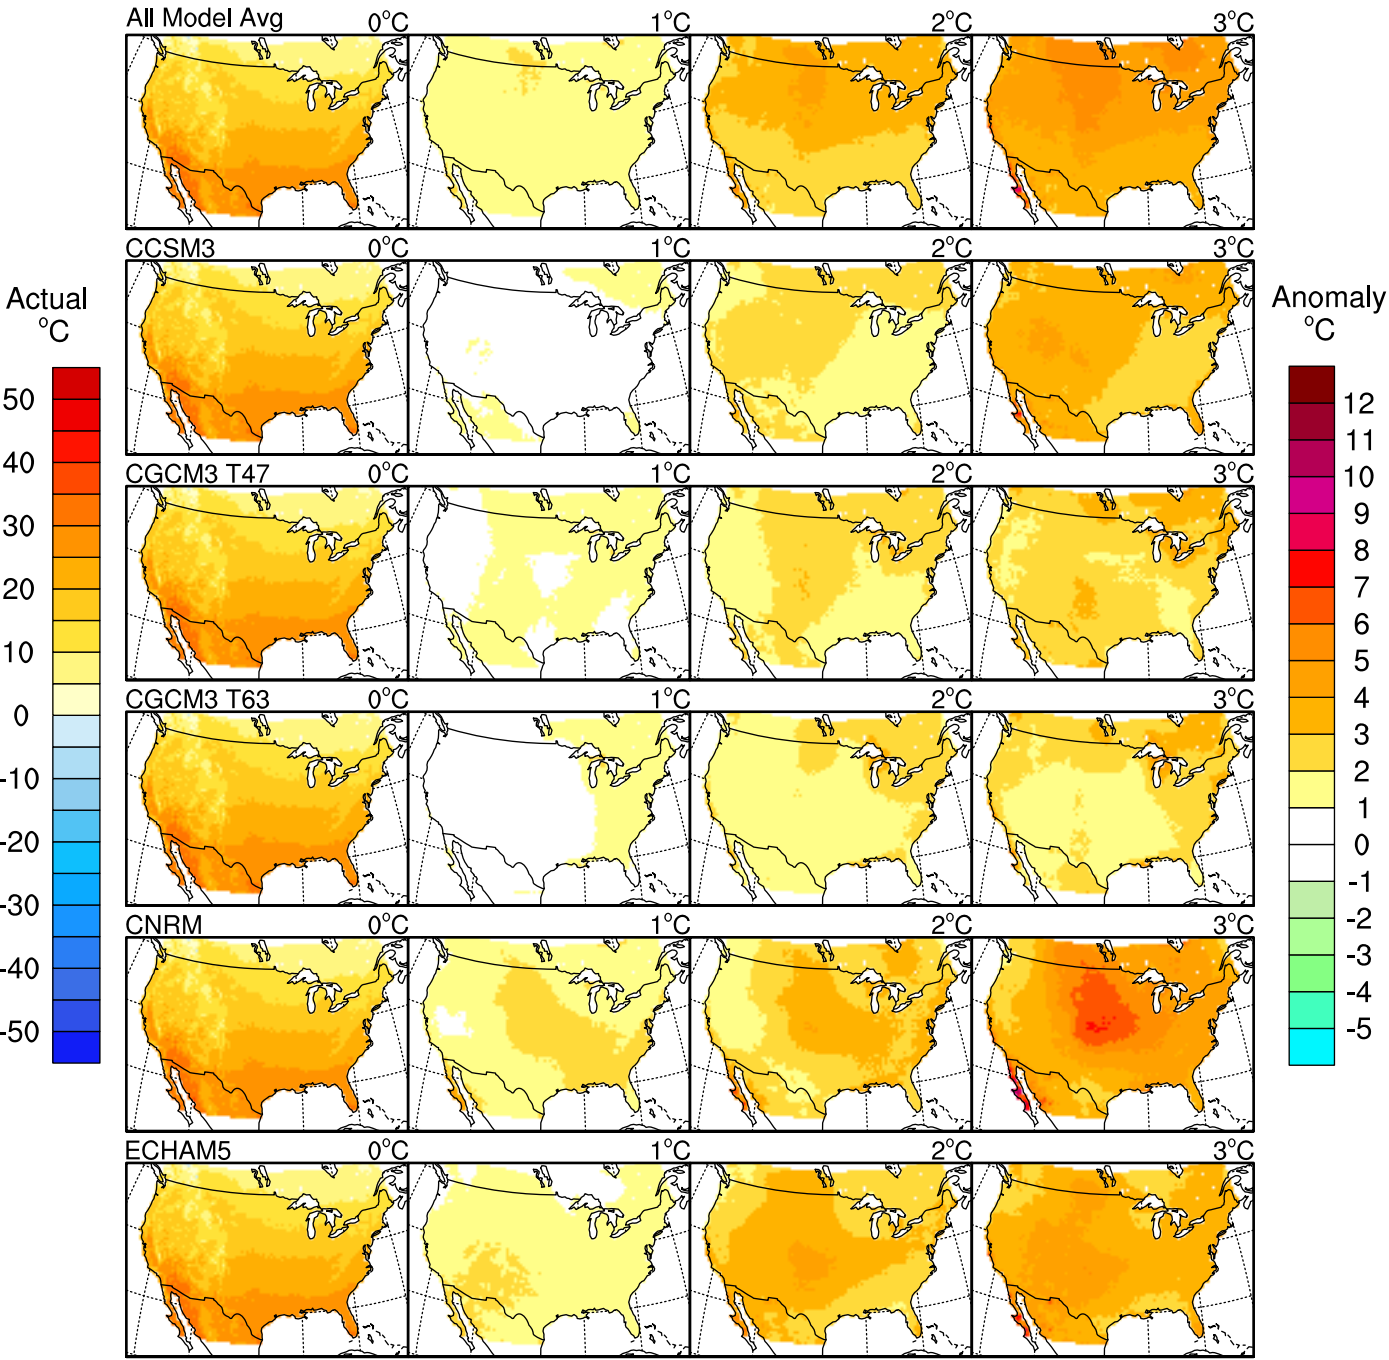
Avg Summer (JJA) Tmax



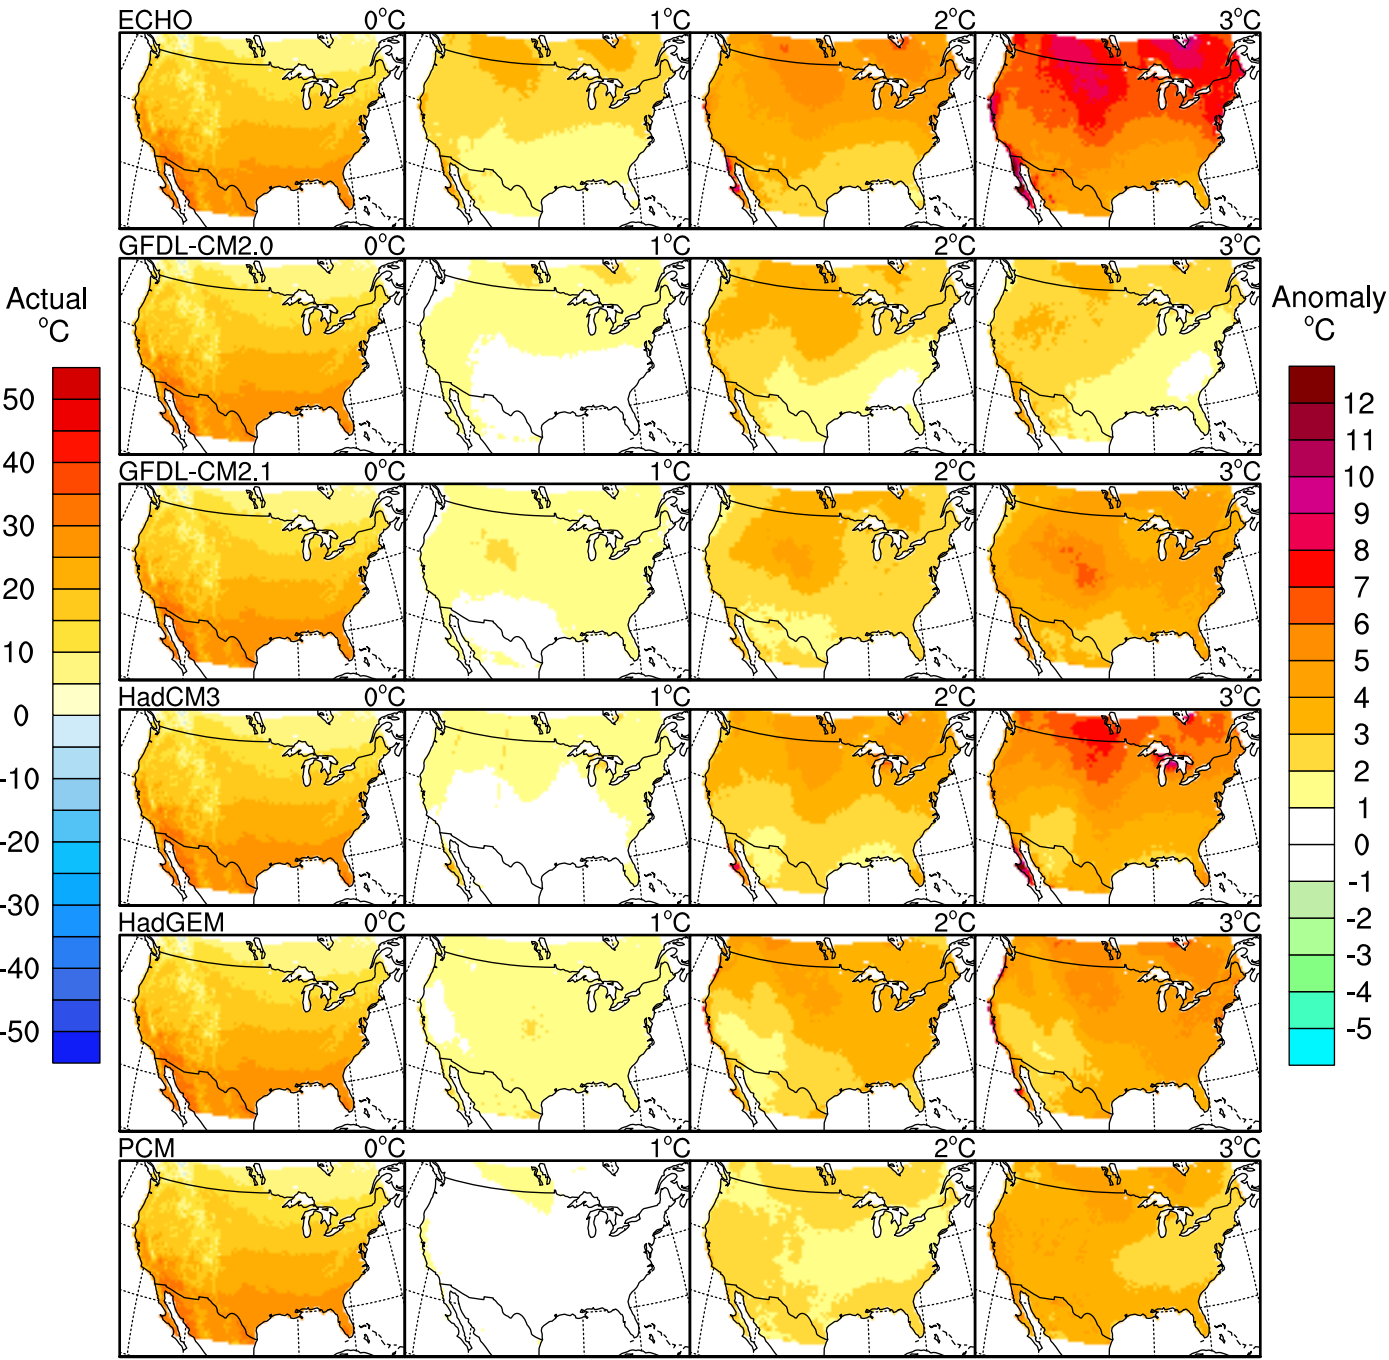
Avg Summer (JJA) Tmax



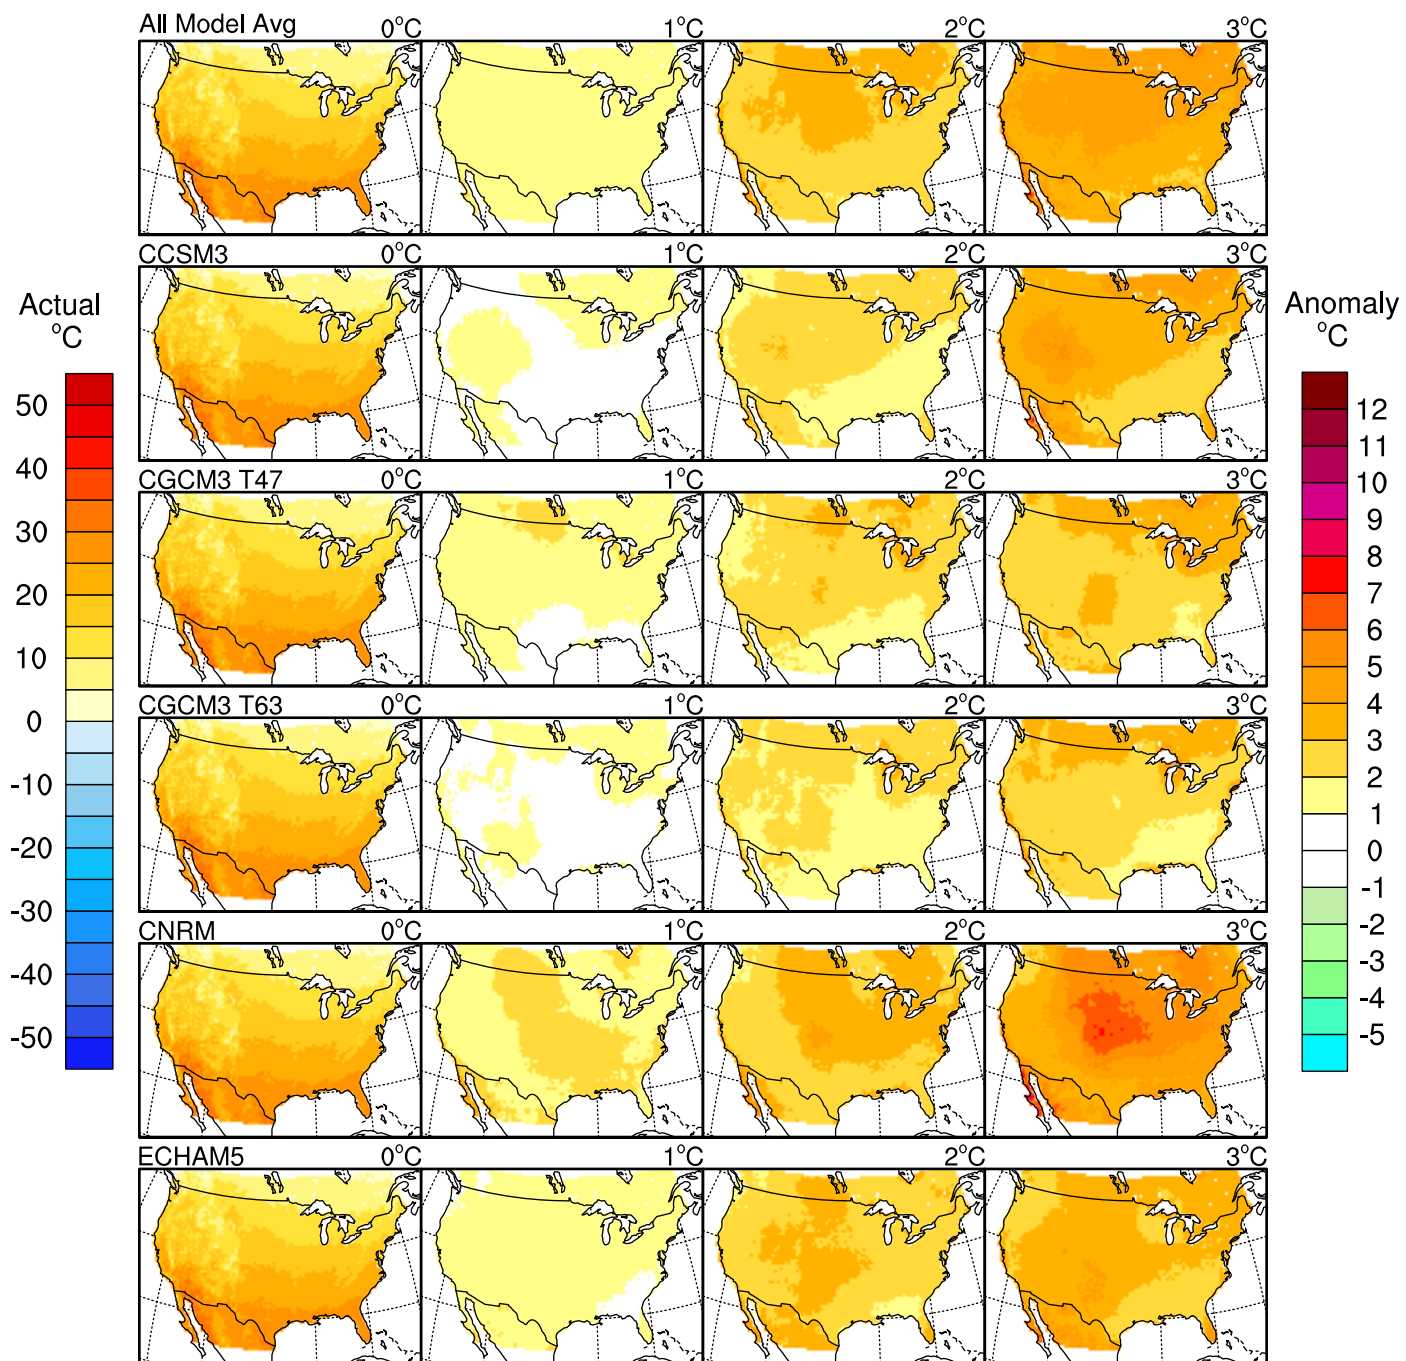
Avg Fall (SON) Tmax



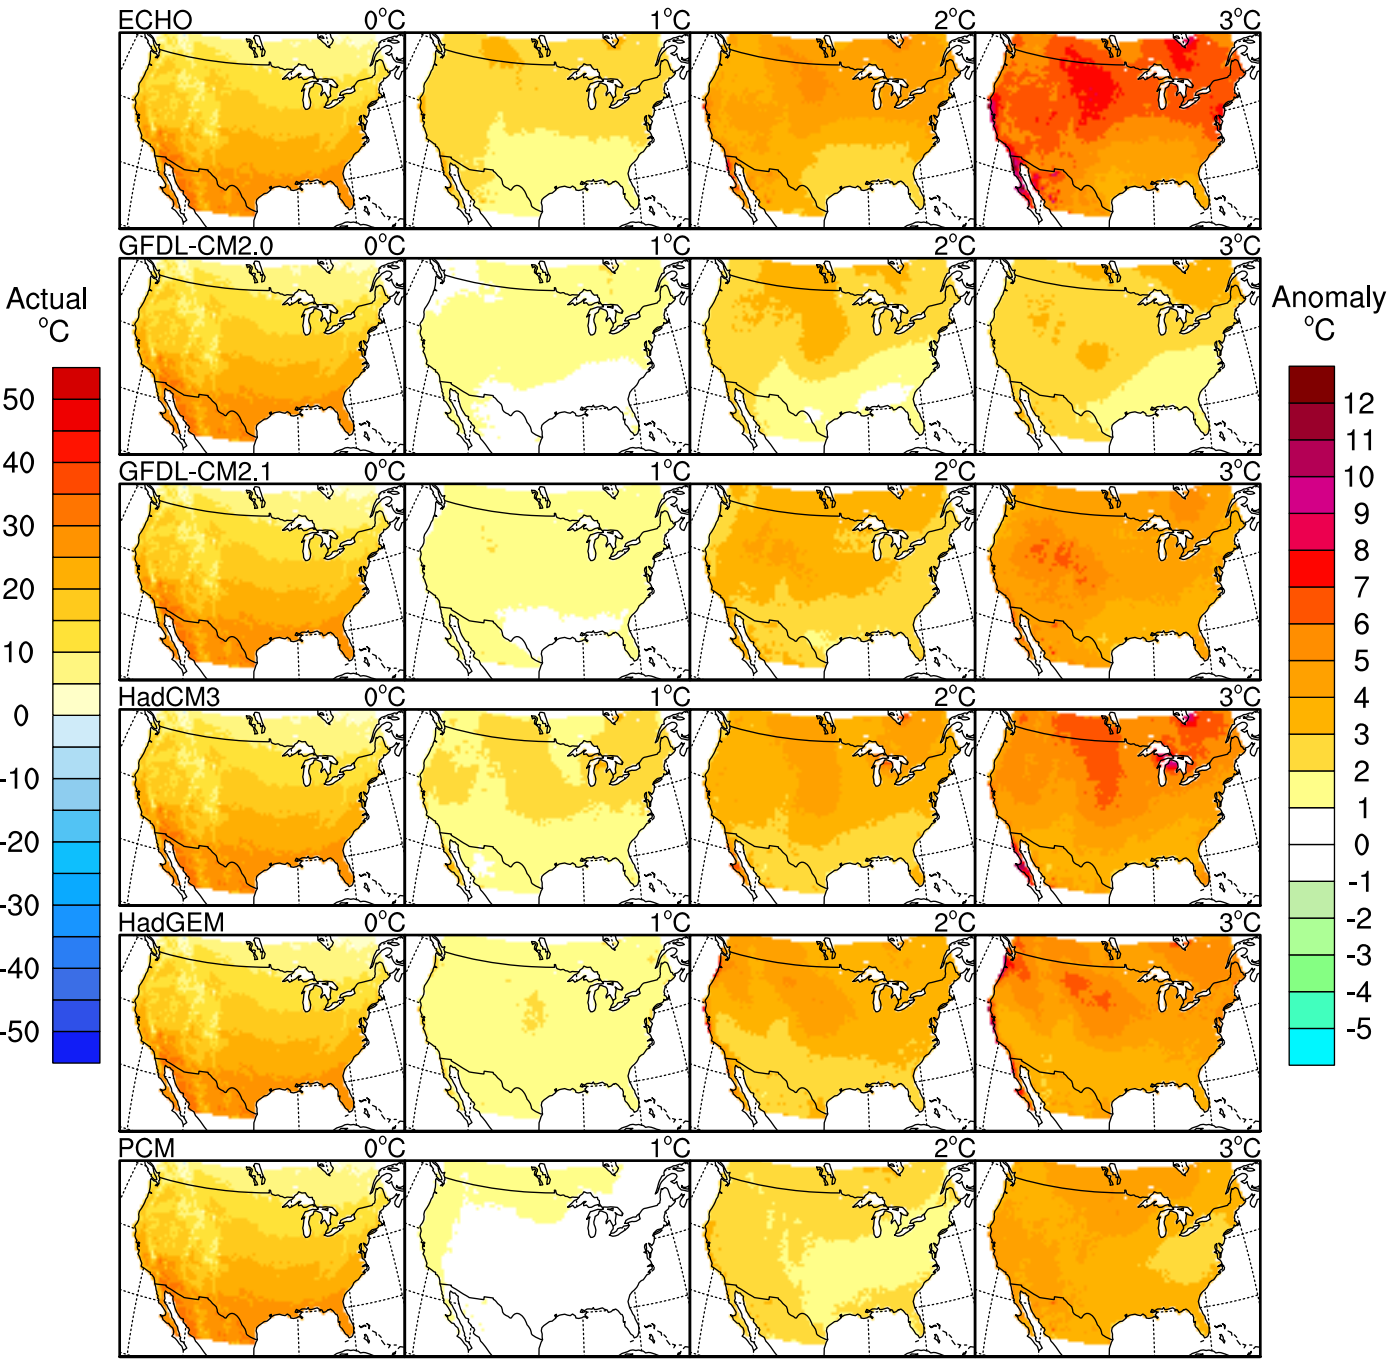
Avg Fall (SON) Tmax




Avg Annual Tmax




Avg Annual Tmax



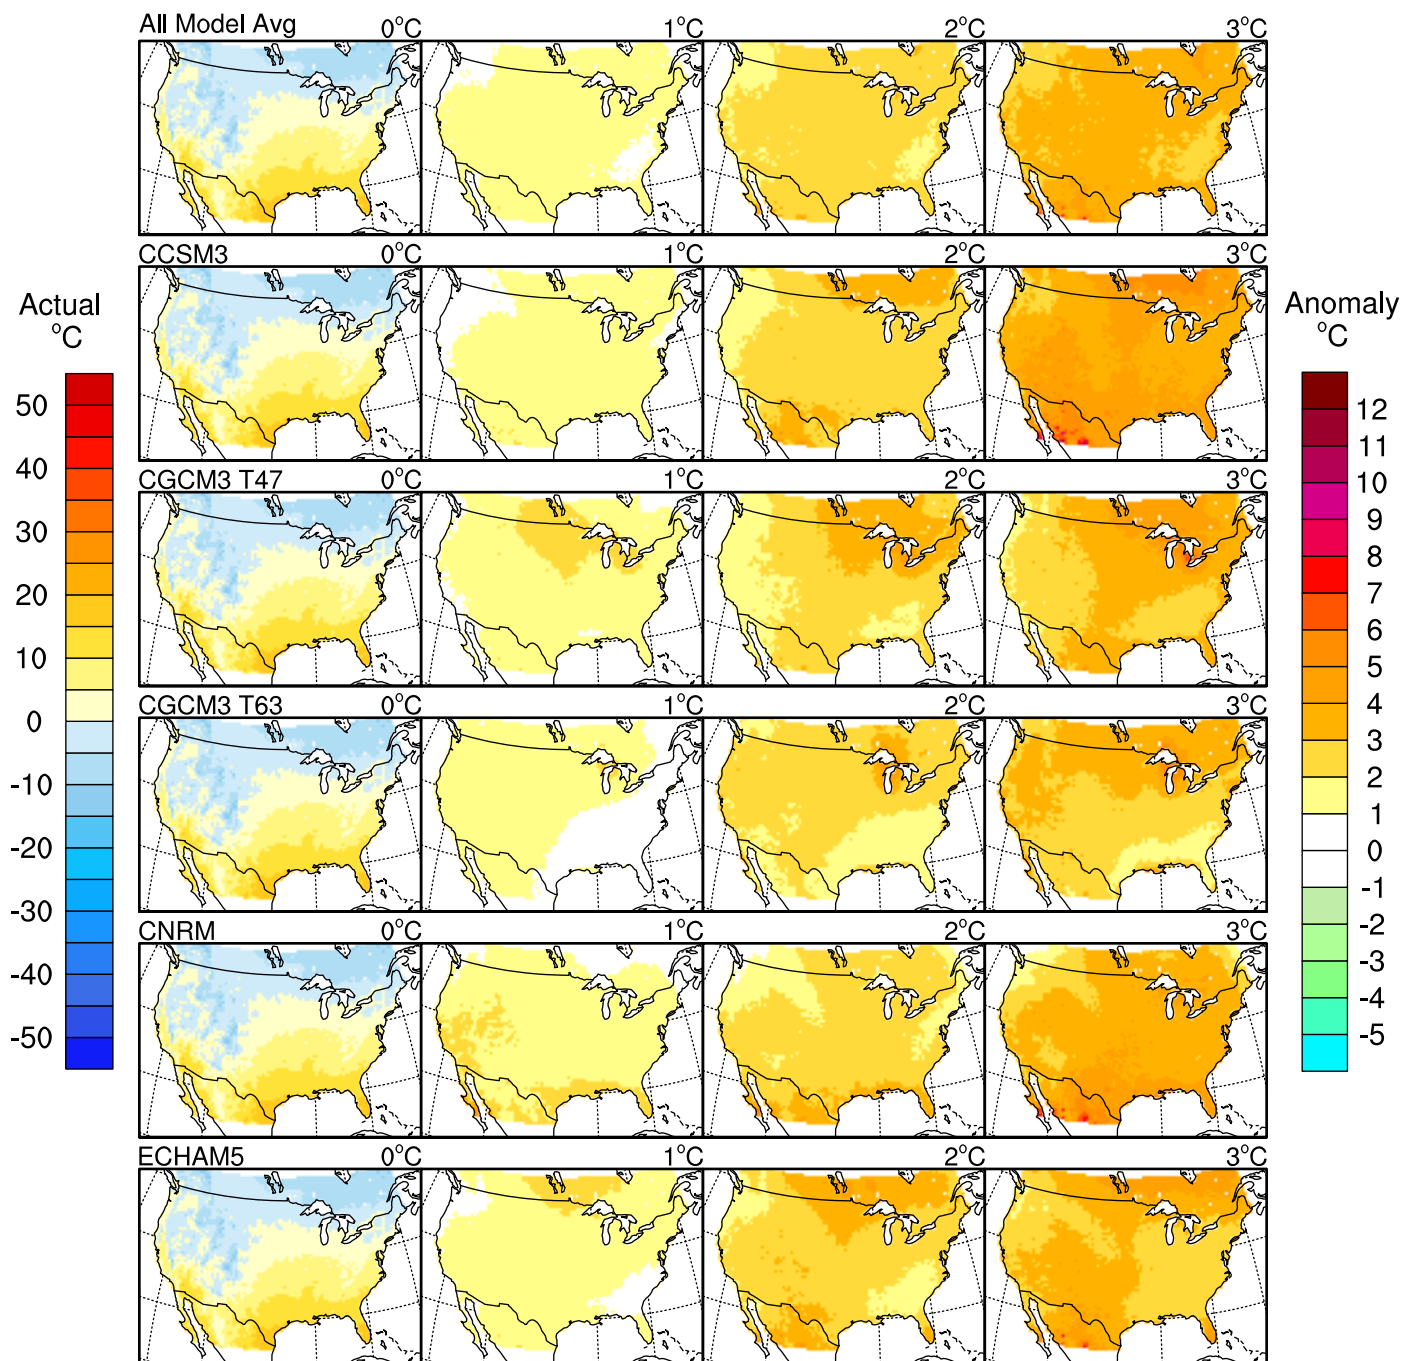
Avg Winter (DJF) Tmin



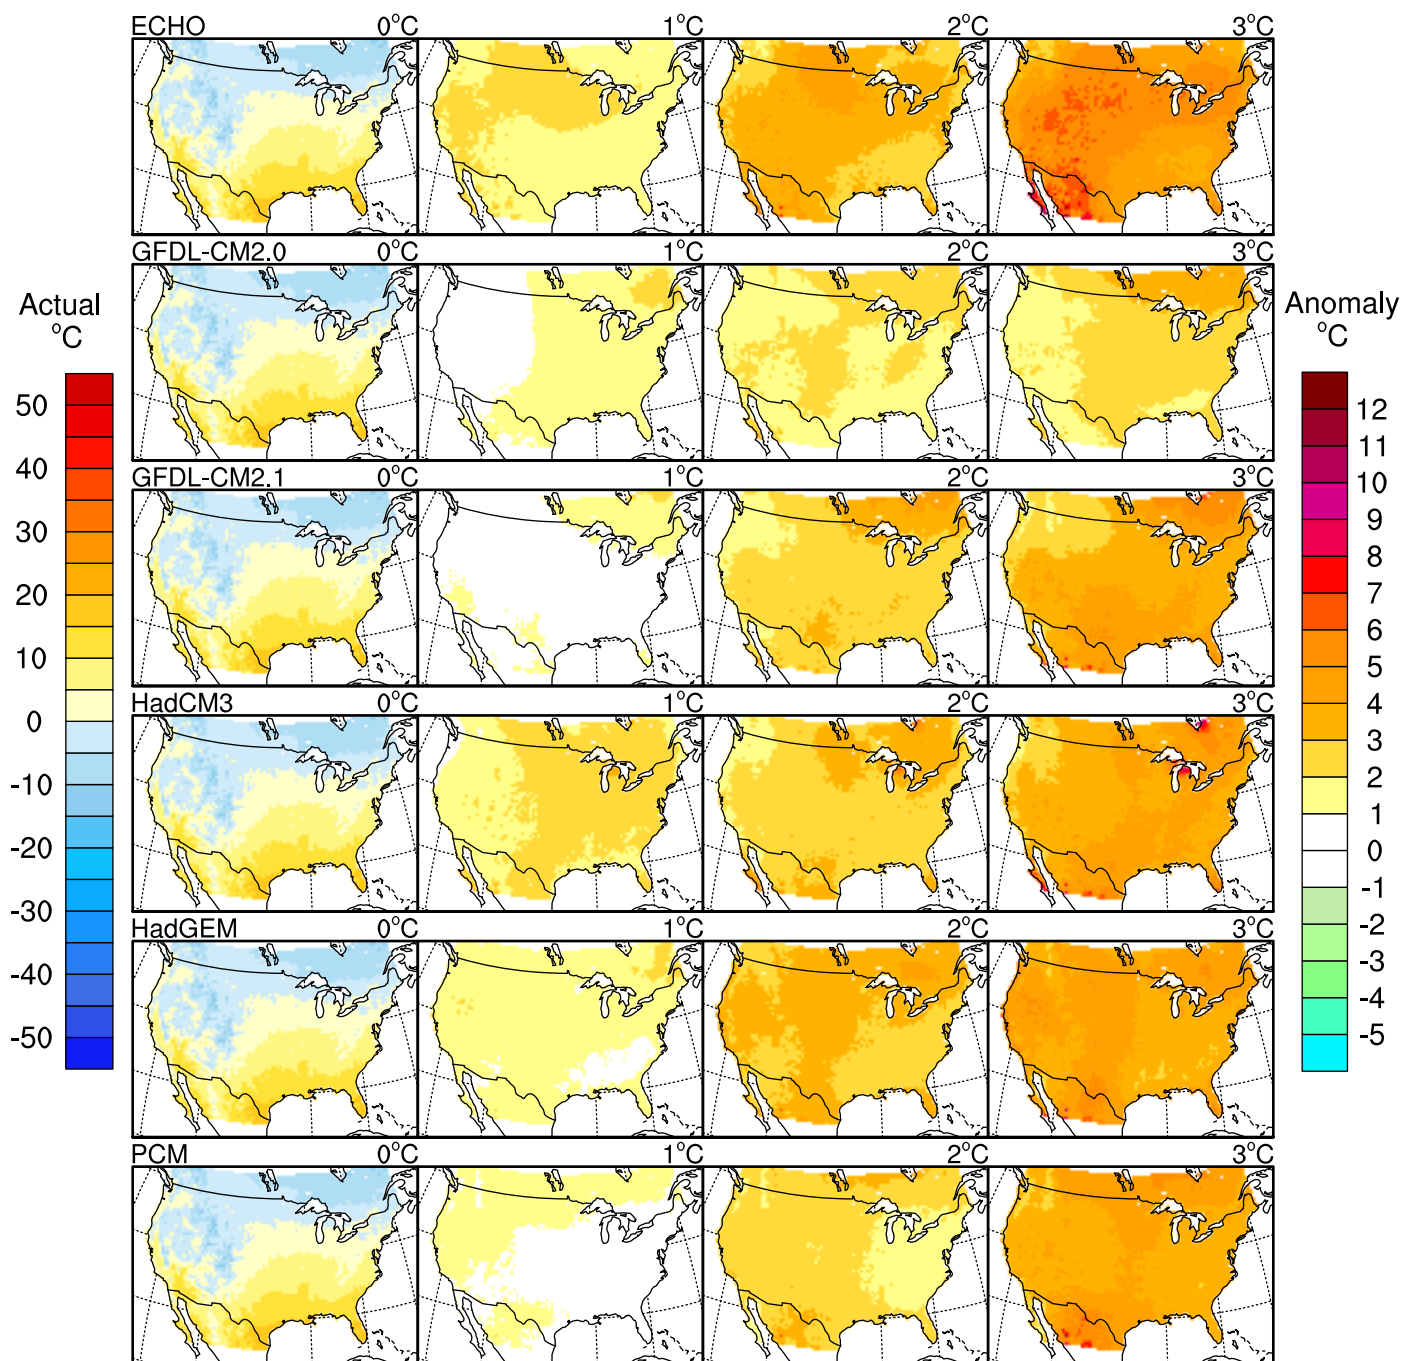
Avg Winter (DJF) Tmin



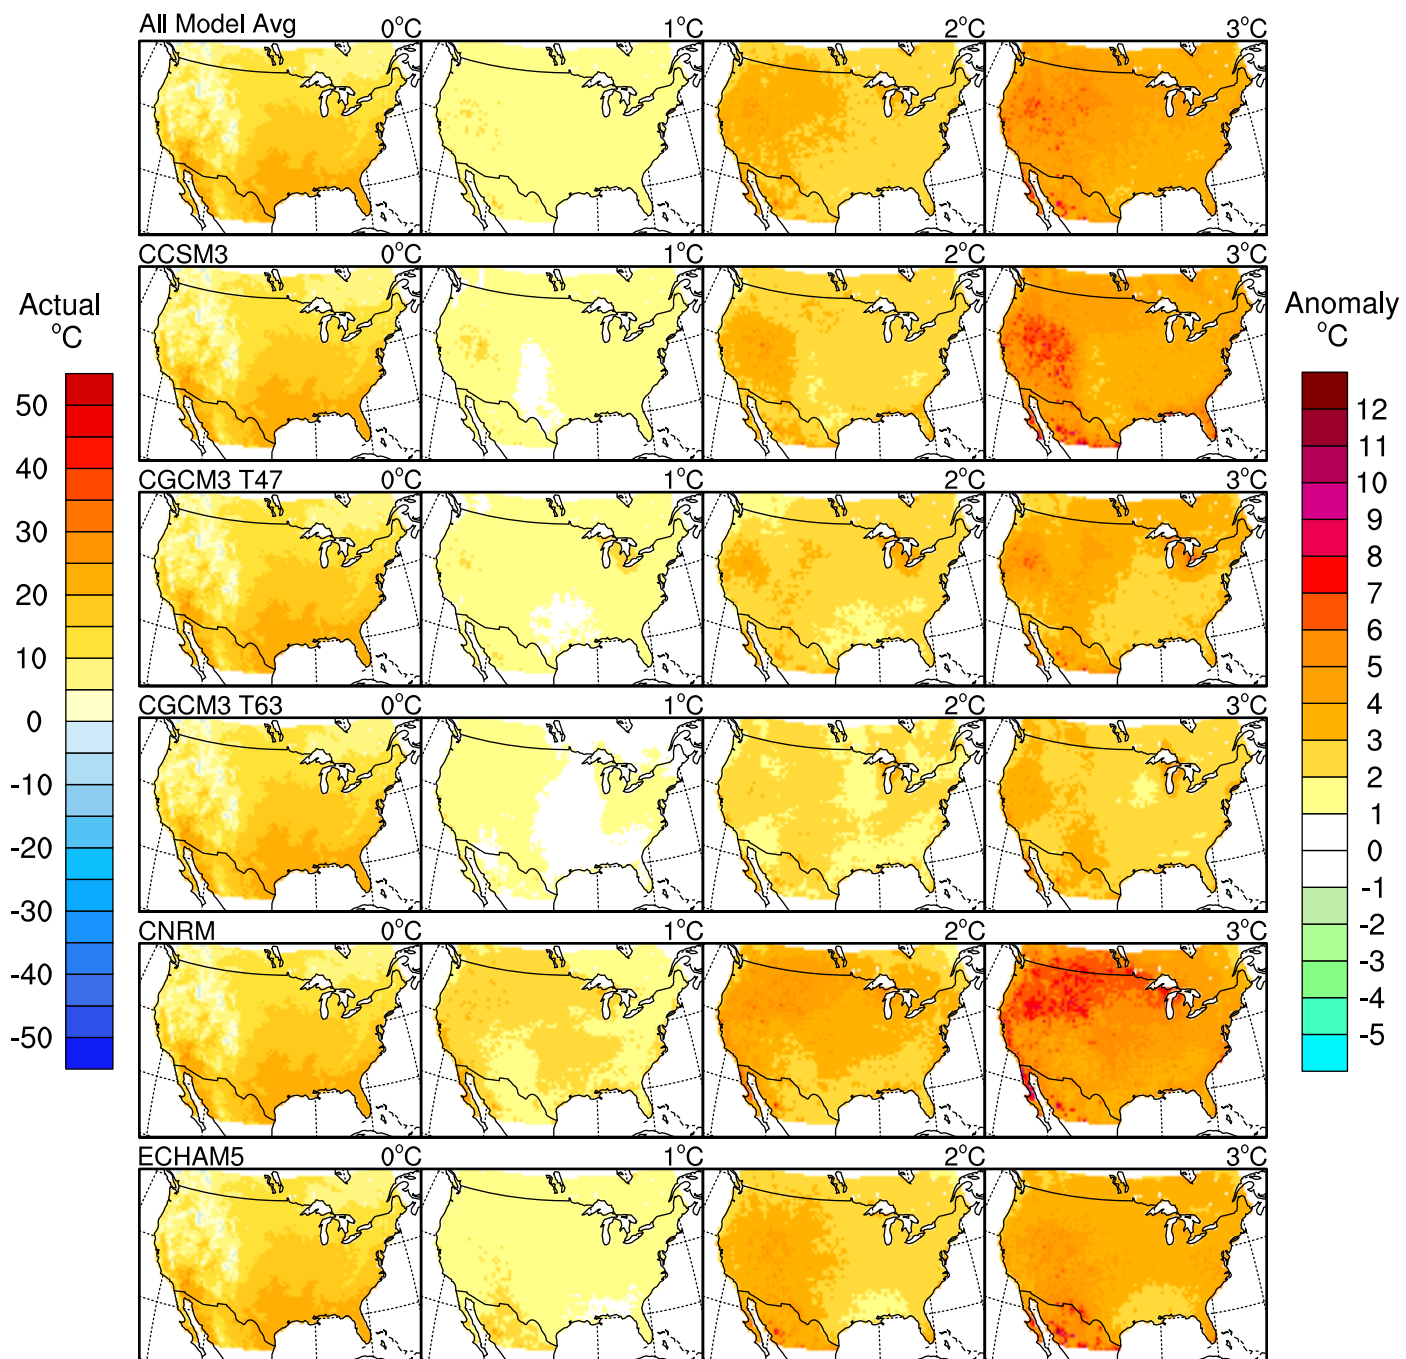
Avg Spring (MAM) Tmin



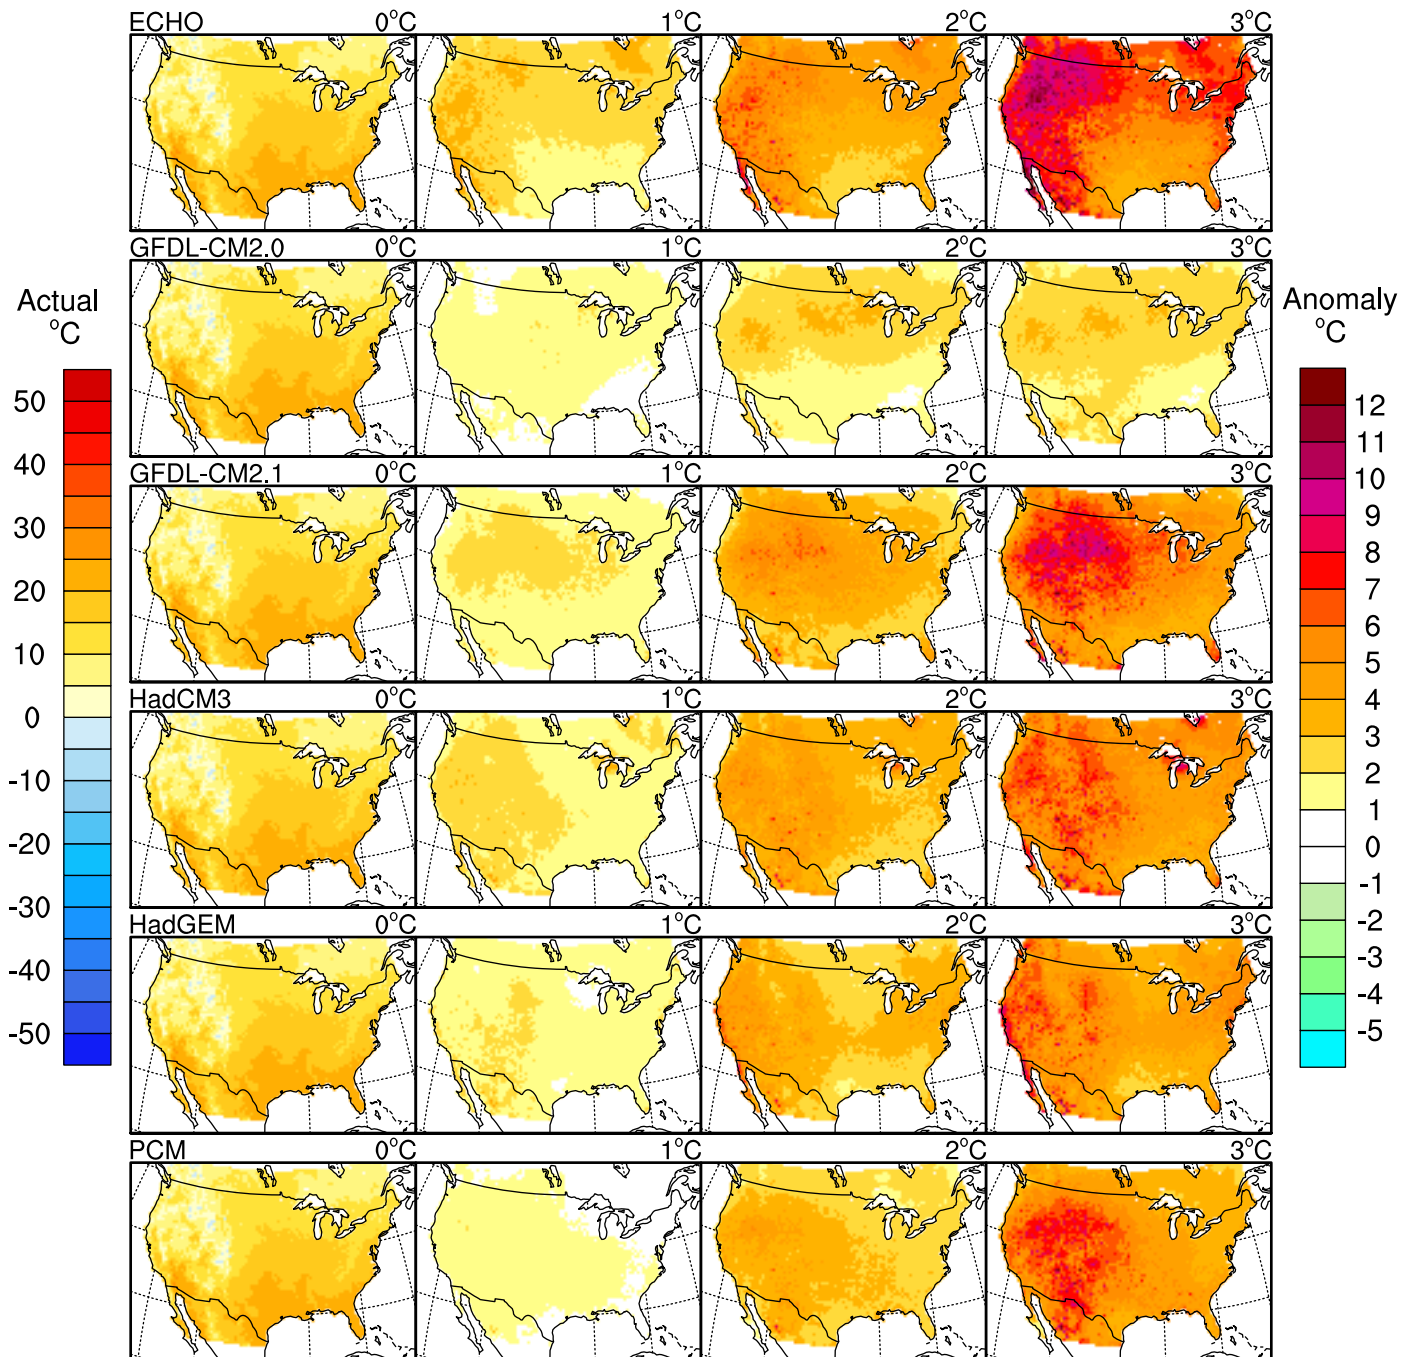
Avg Spring (MAM) Tmin



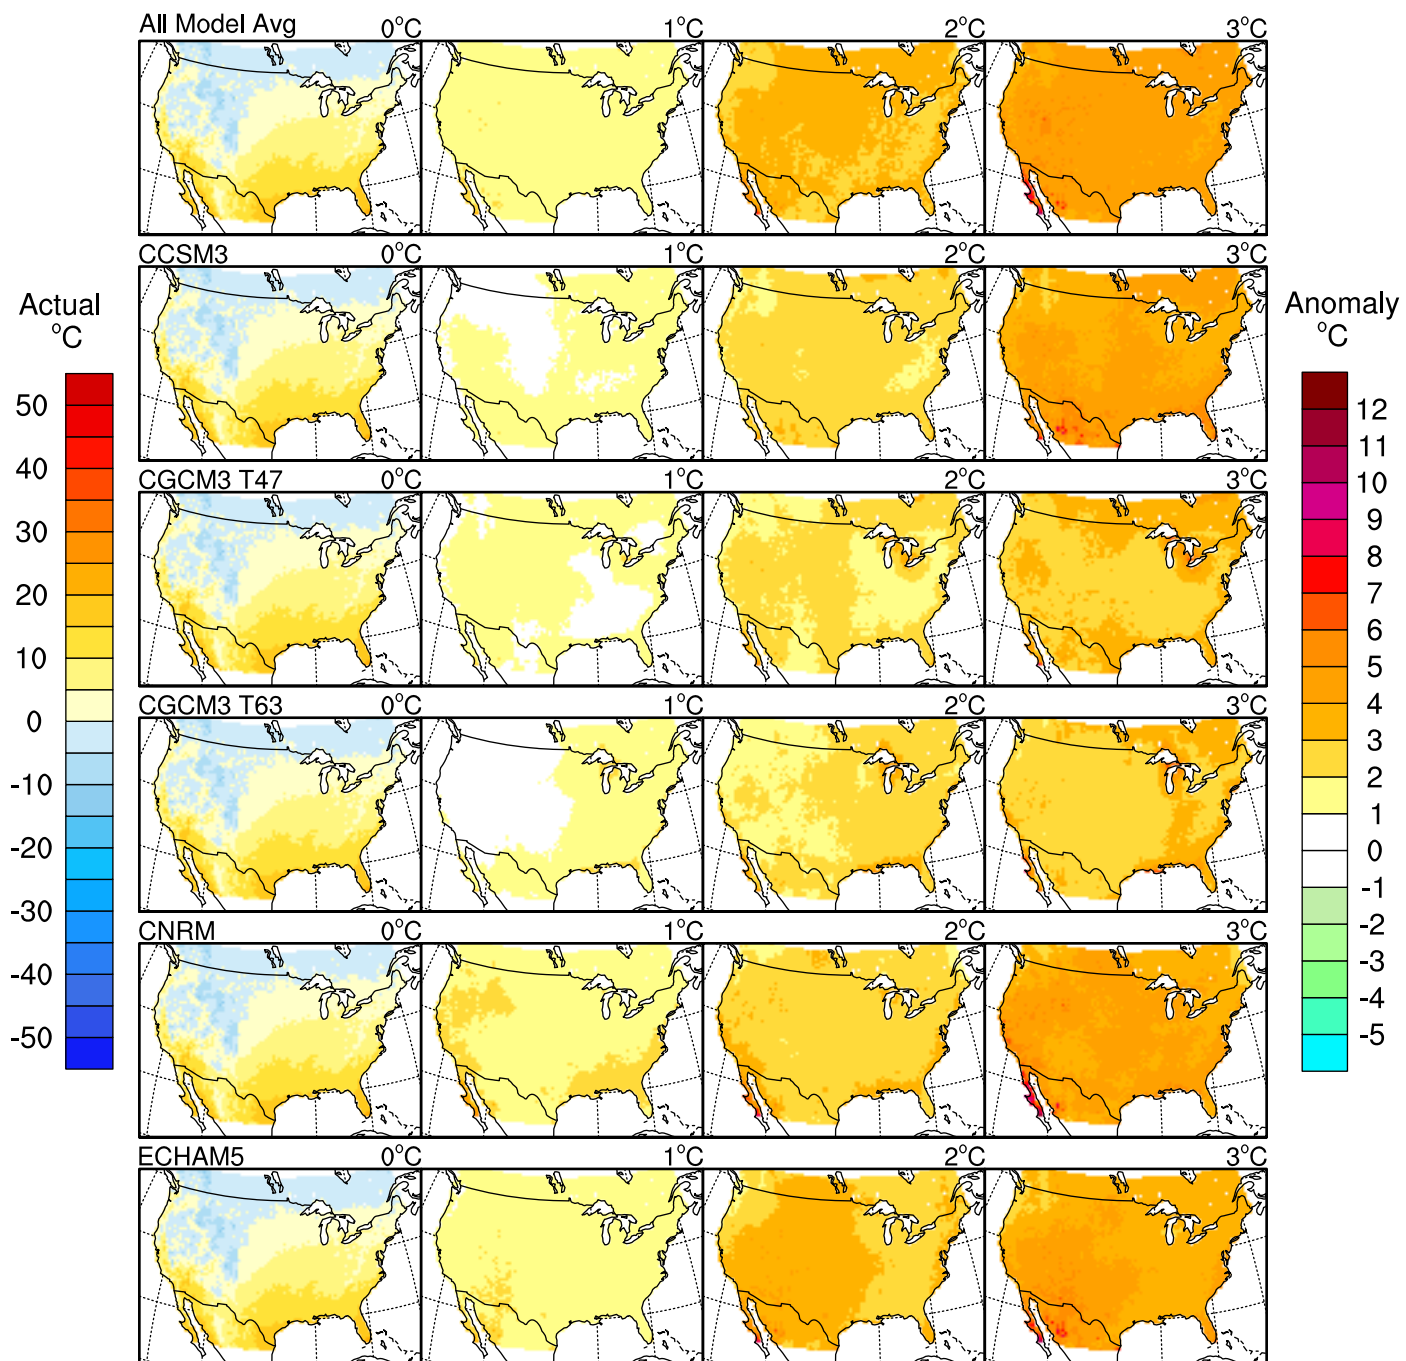
Avg Summer (JJA) Tmin




Avg Summer (JJA) Tmin




Avg Fall (SON) Tmin




Avg Fall (SON) Tmin



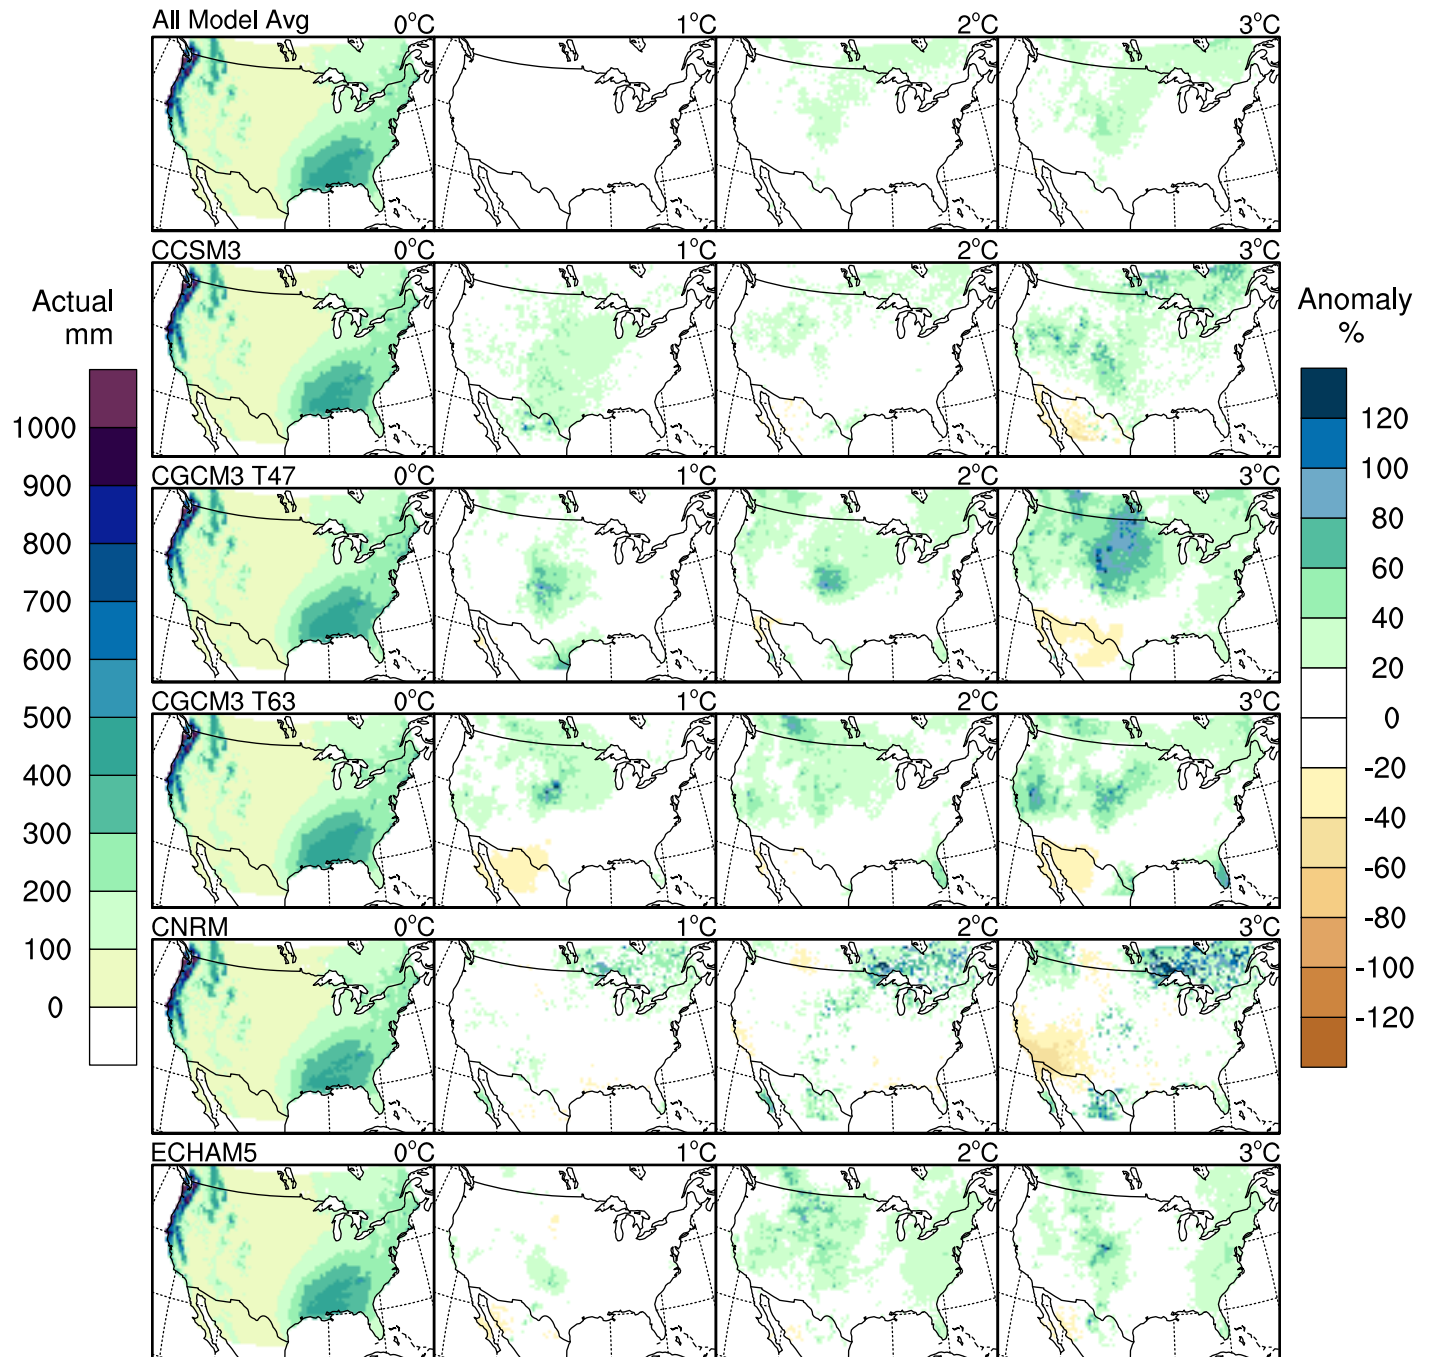
Avg Annual Tmin



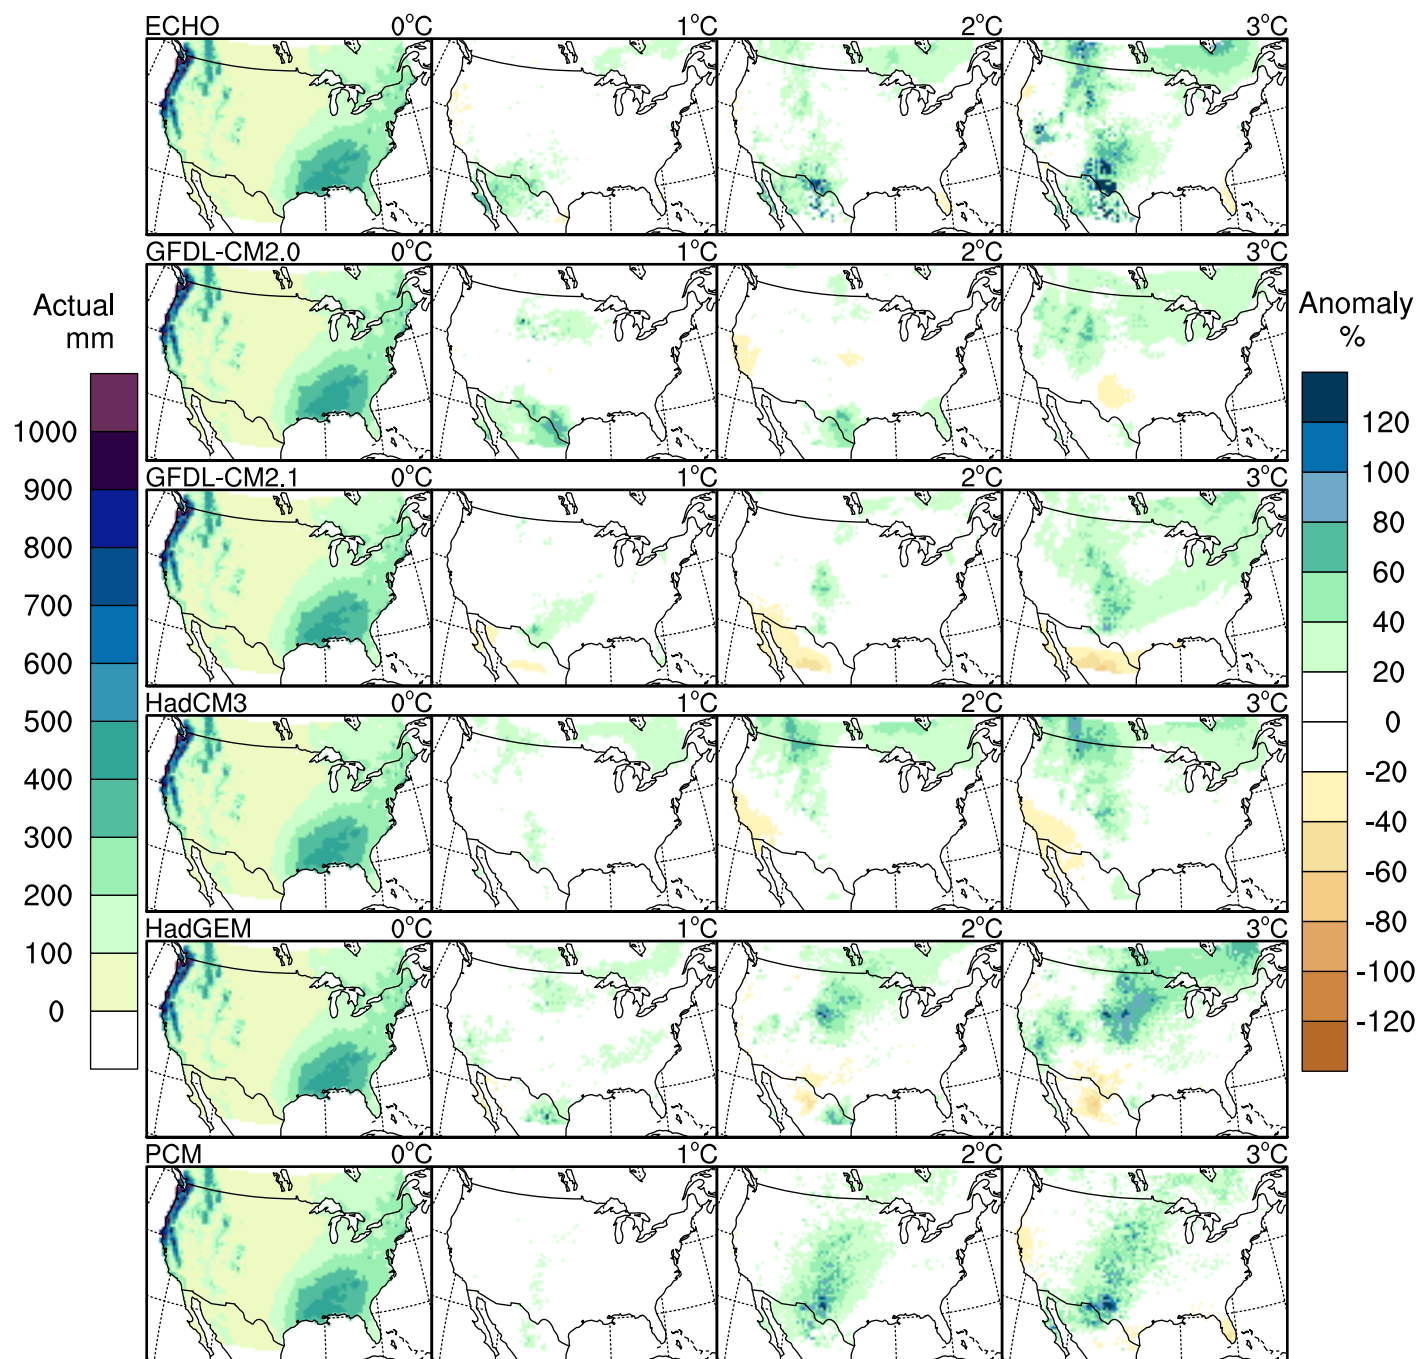
Avg Annual Tmin



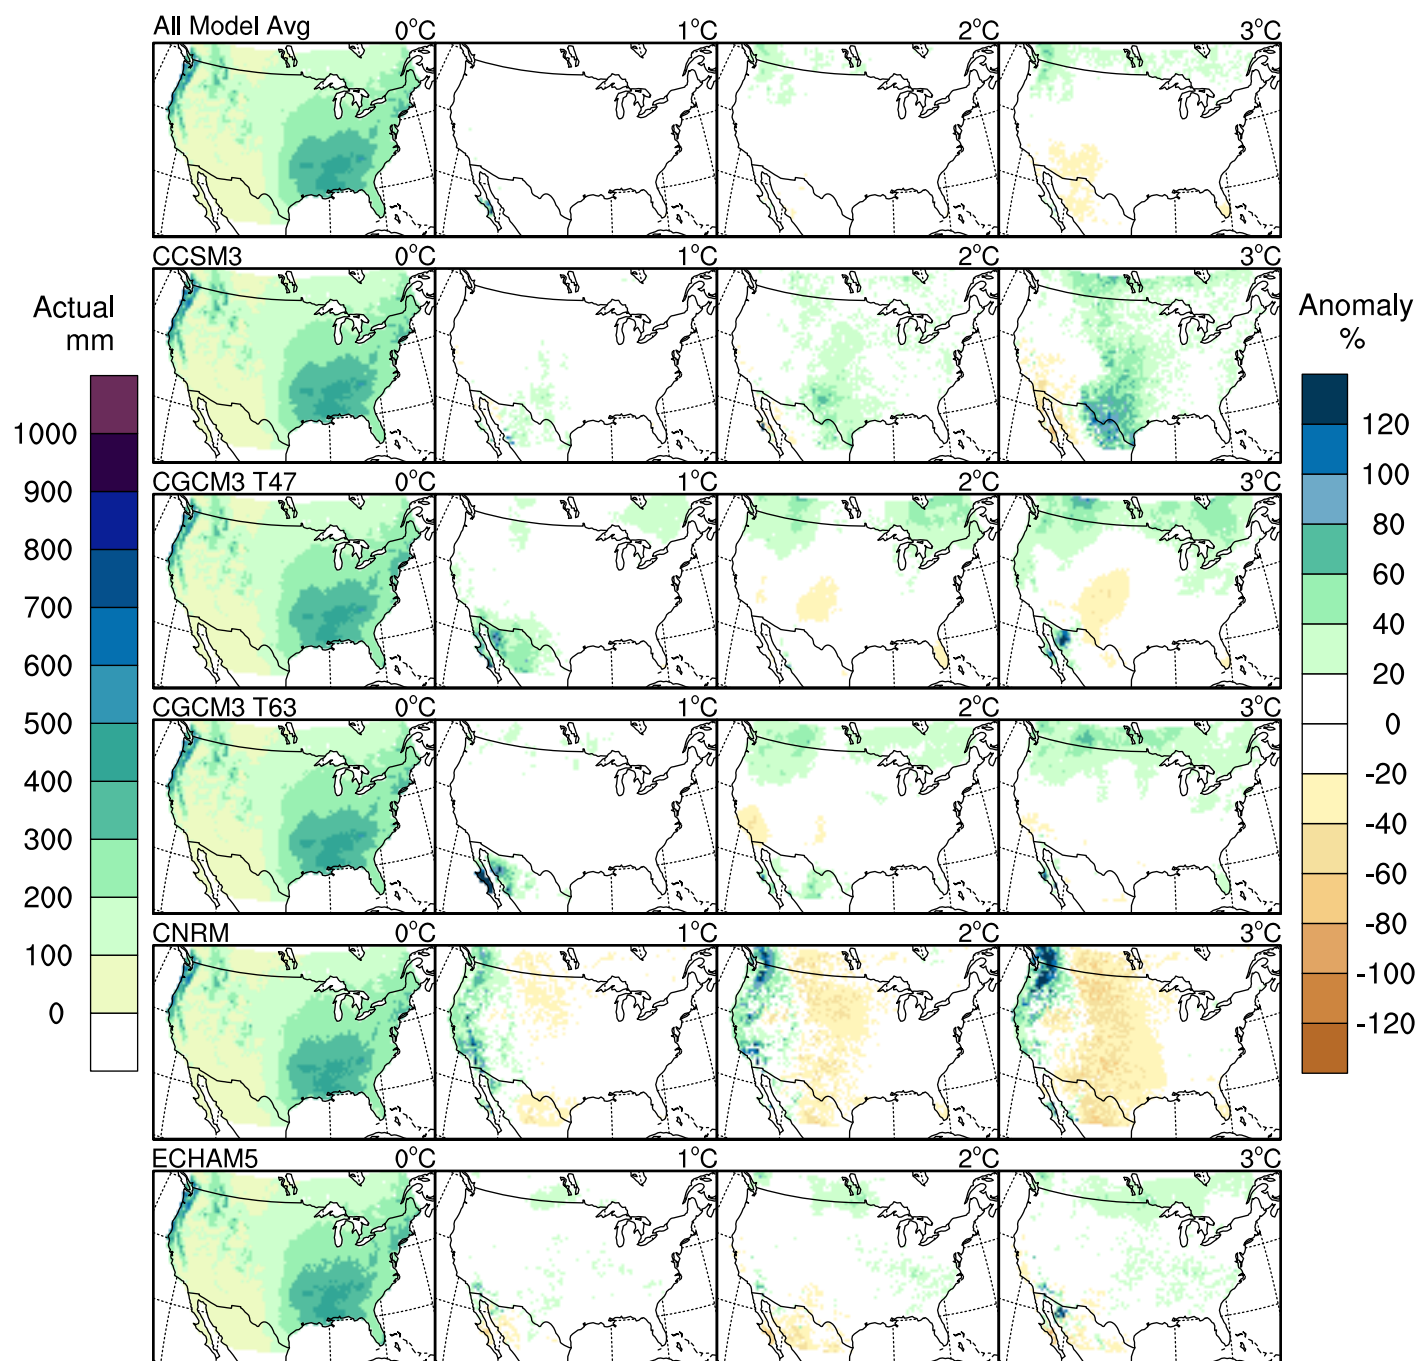
Avg Winter (DJF) Cumulative Pr



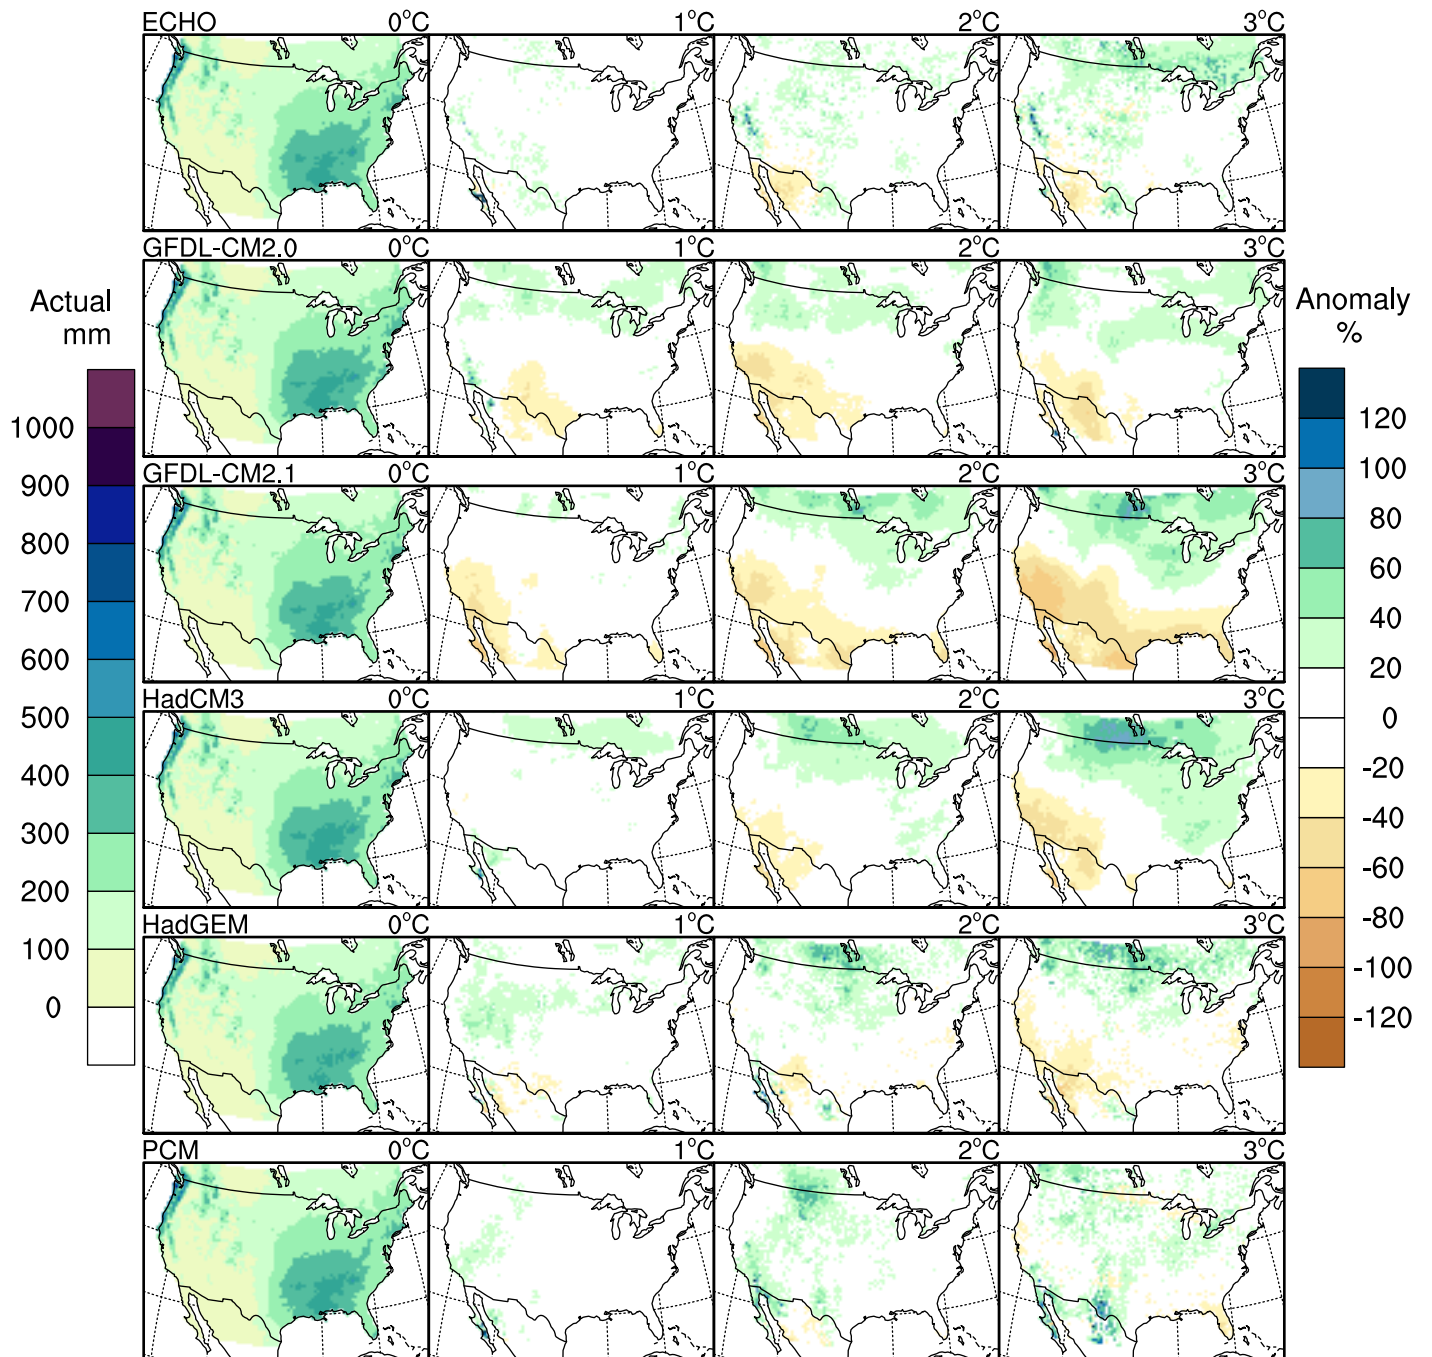
Avg Winter (DJF) Cumulative Pr



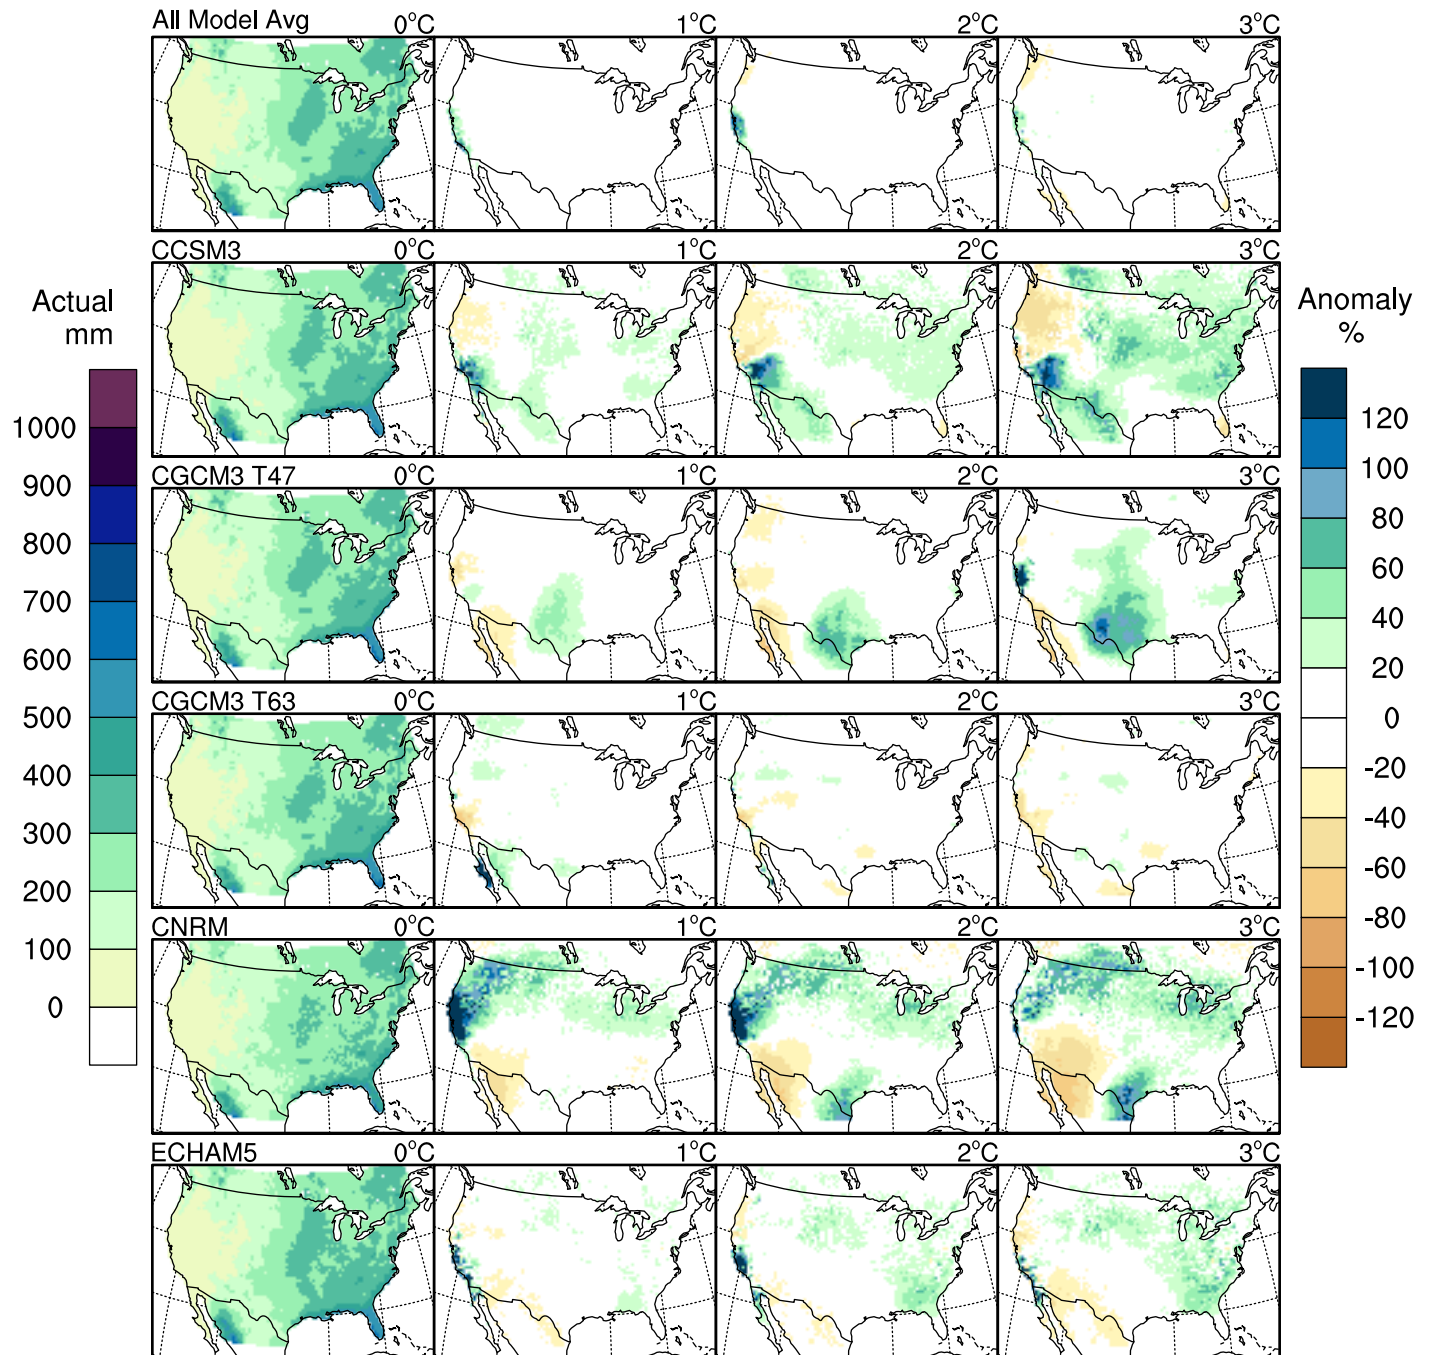
Avg Spring (MAM) Cumulative Pr



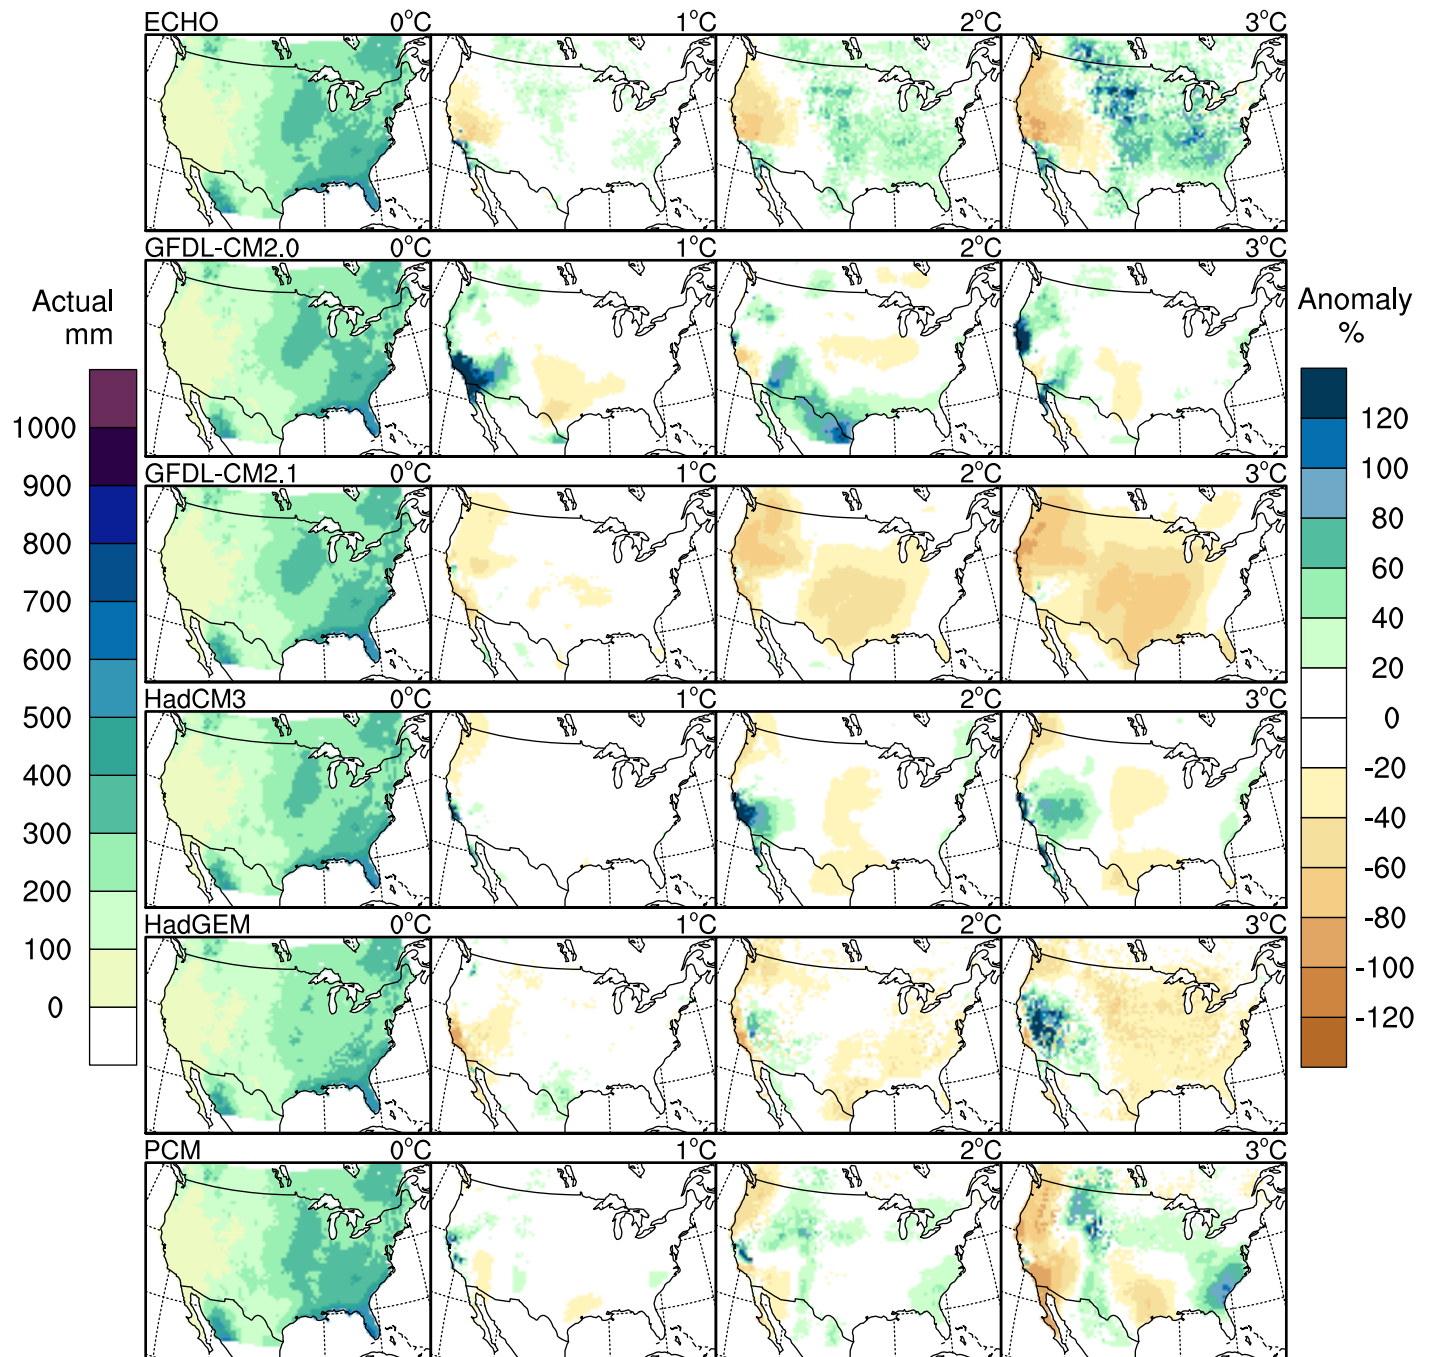
Avg Spring (MAM) Cumulative Pr



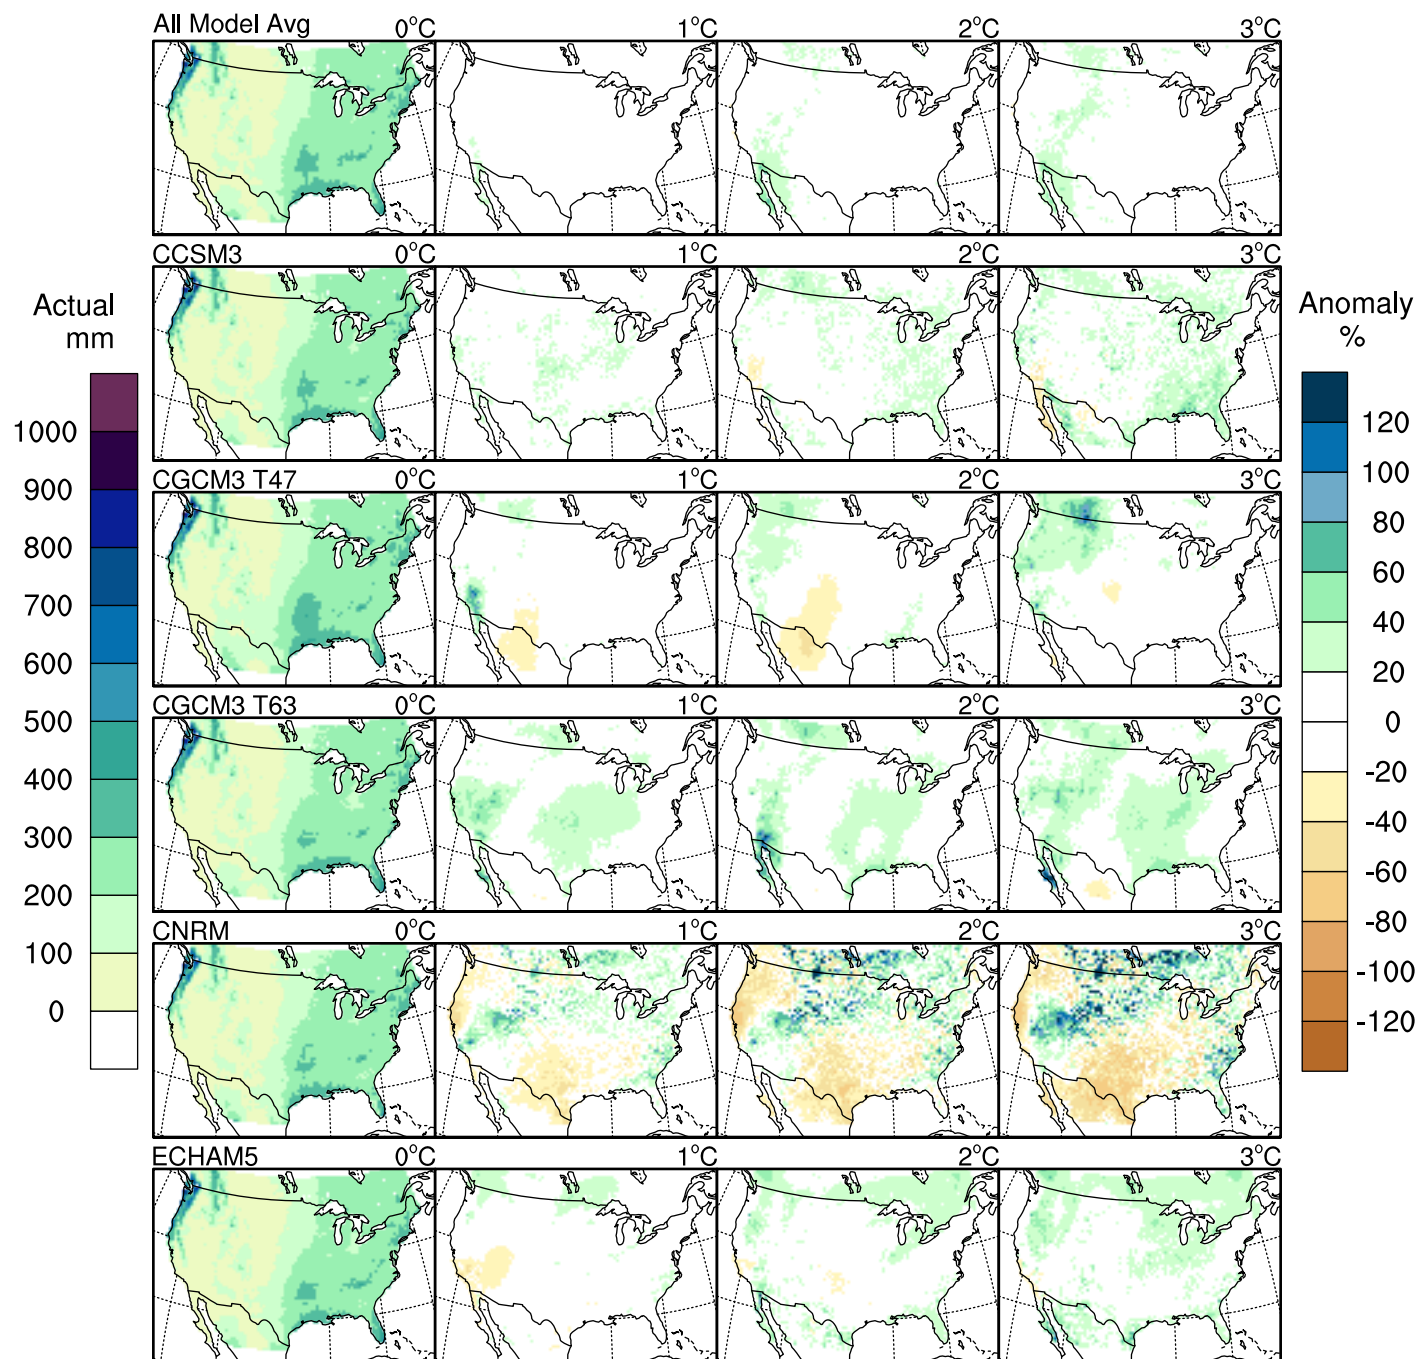
Avg Summer (JJA) Cumulative Pr



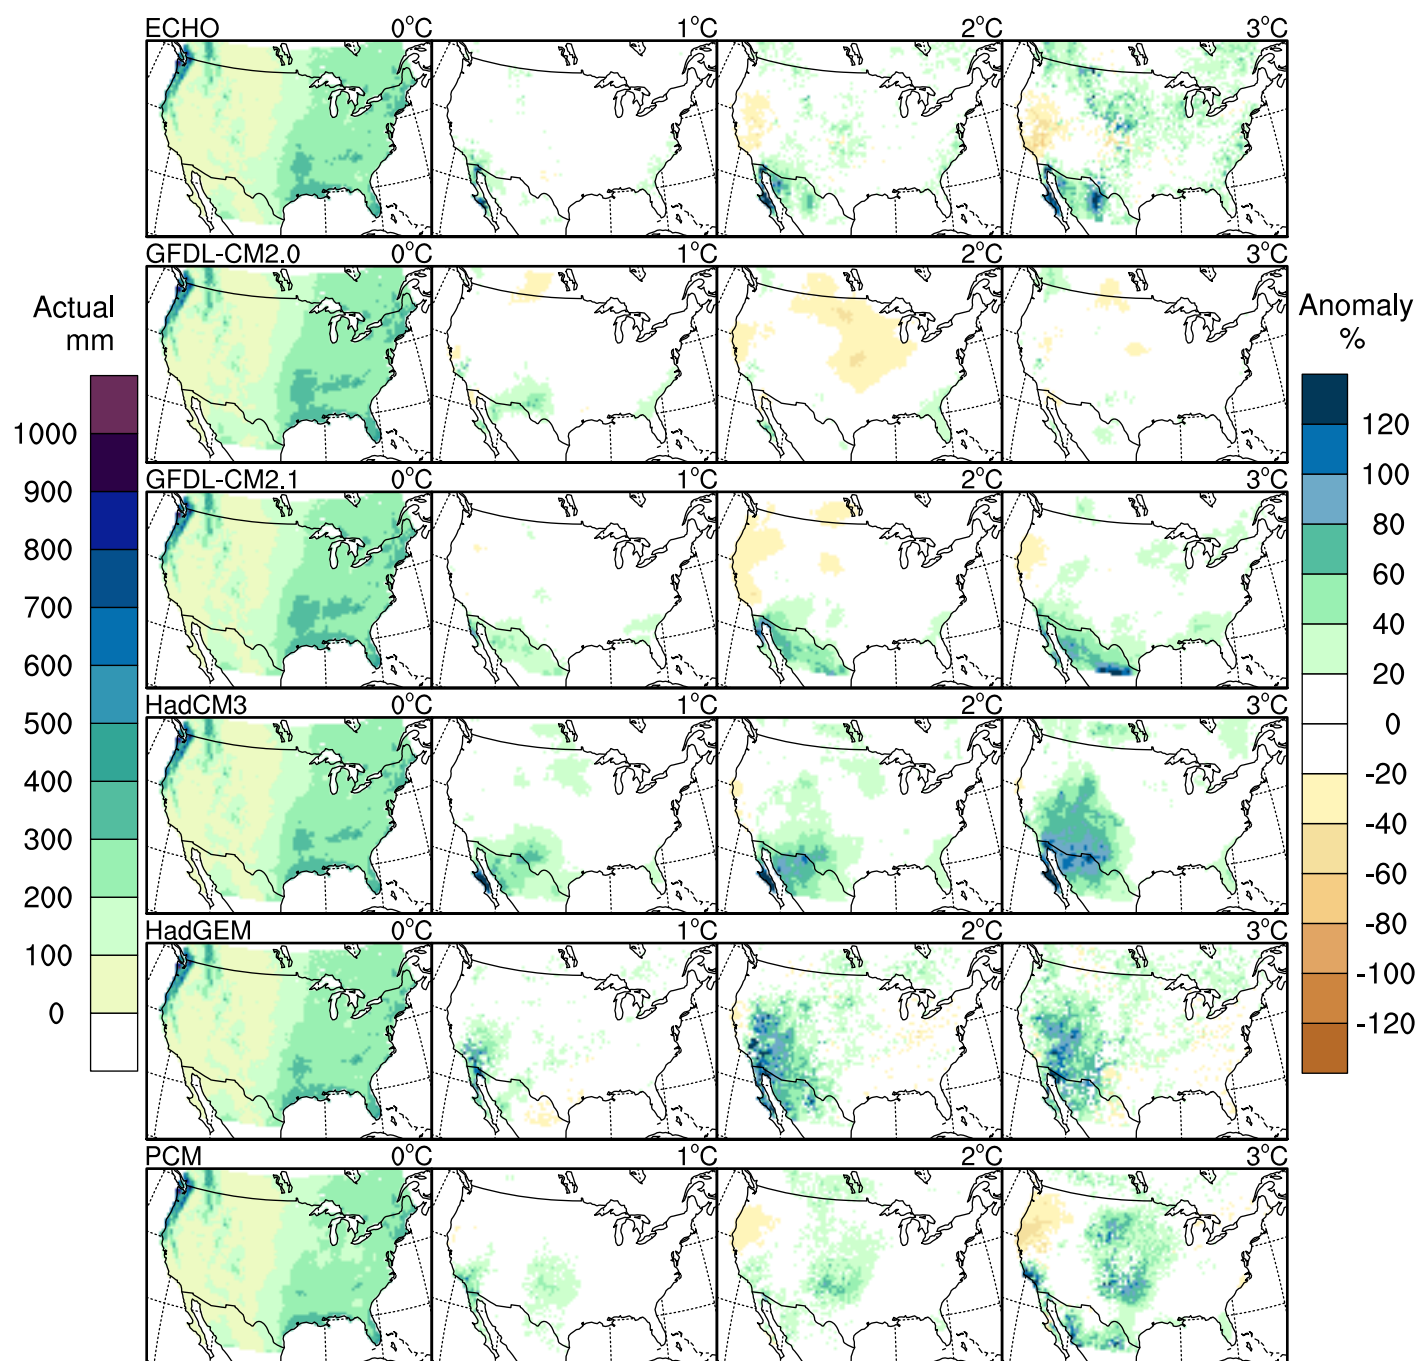
Avg Summer (JJA) Cumulative Pr



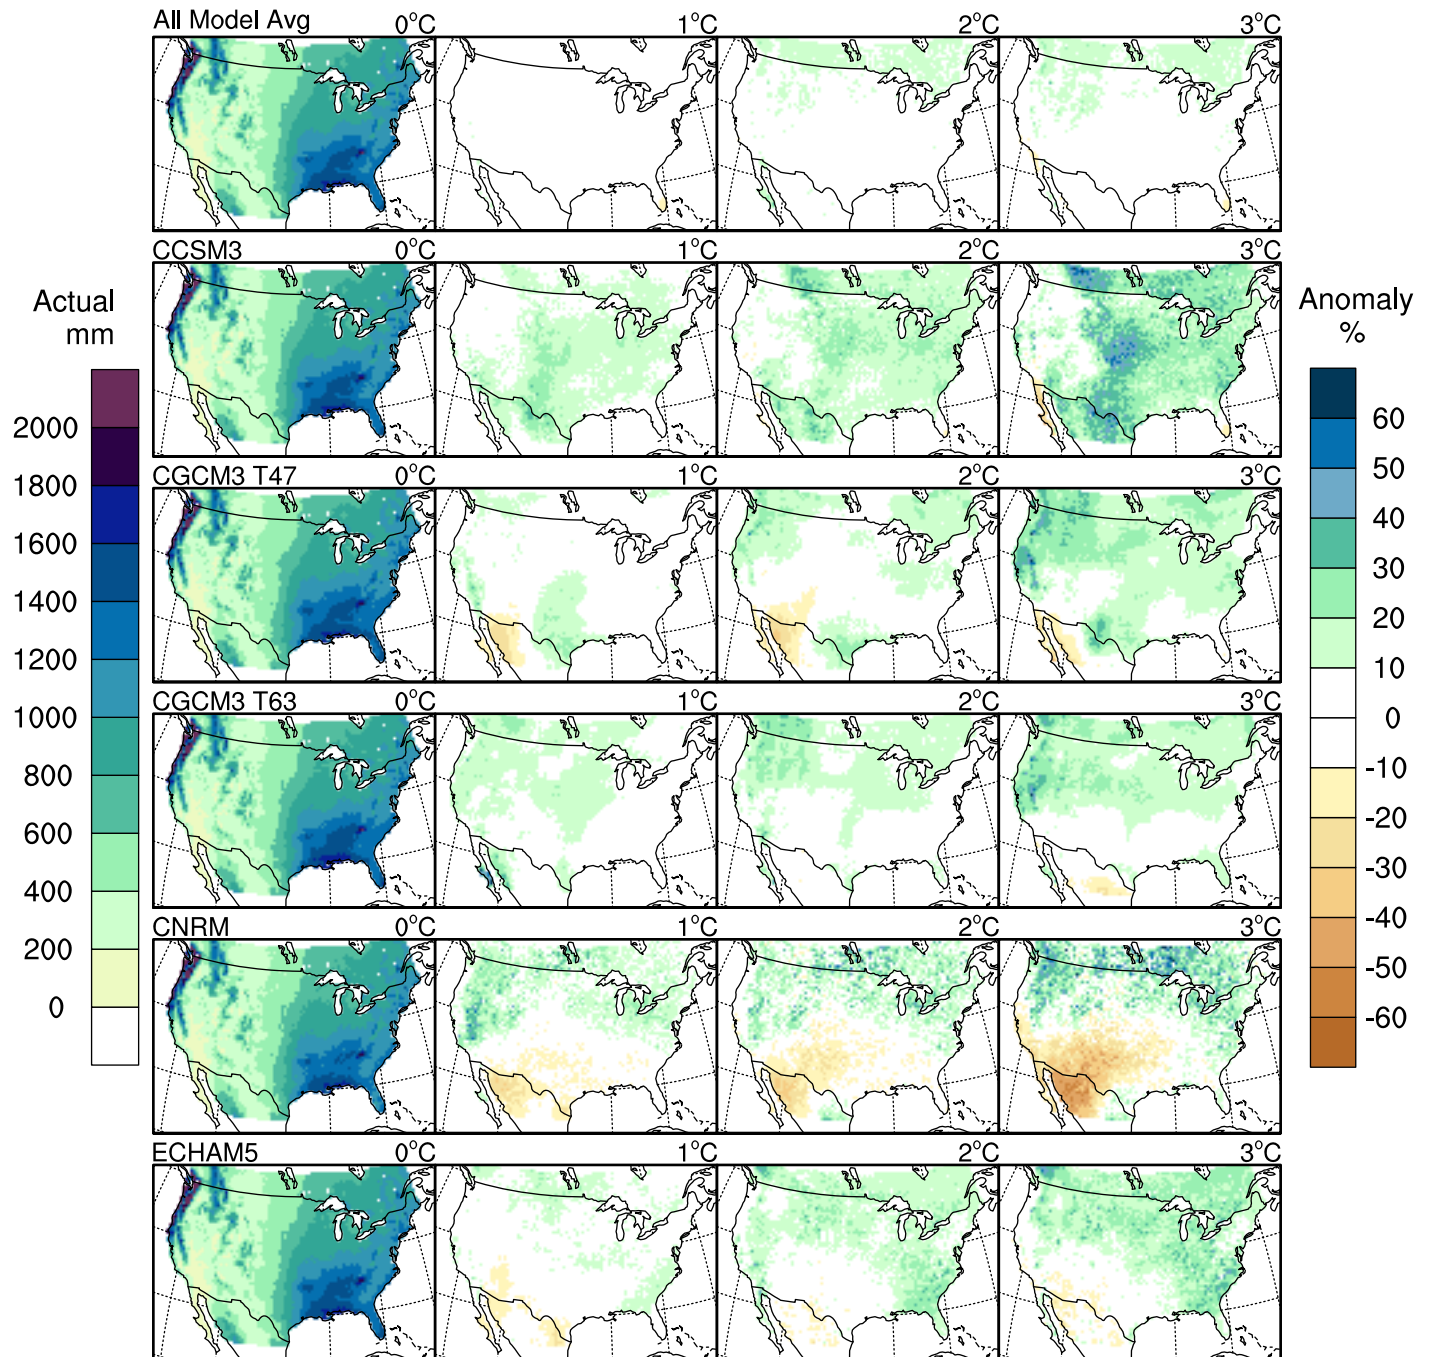
Avg Fall (SON) Cumulative Pr



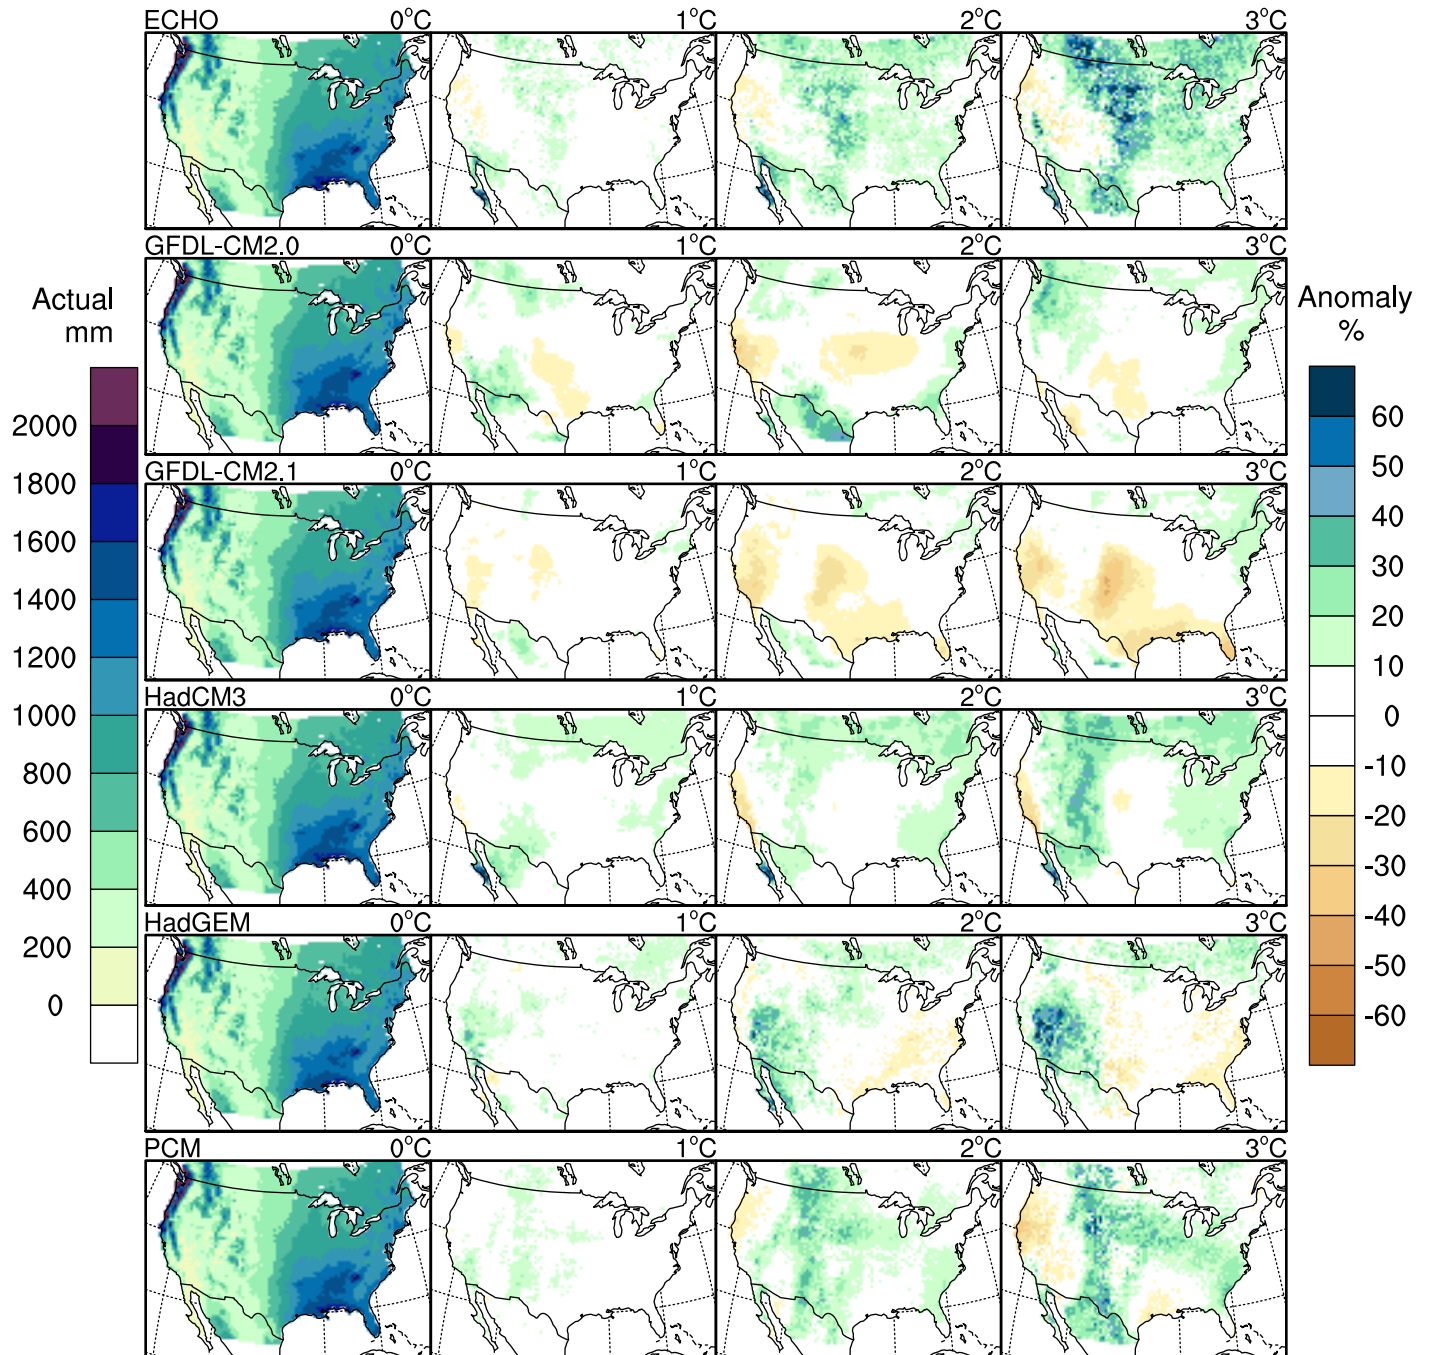
Avg Fall (SON) Cumulative Pr



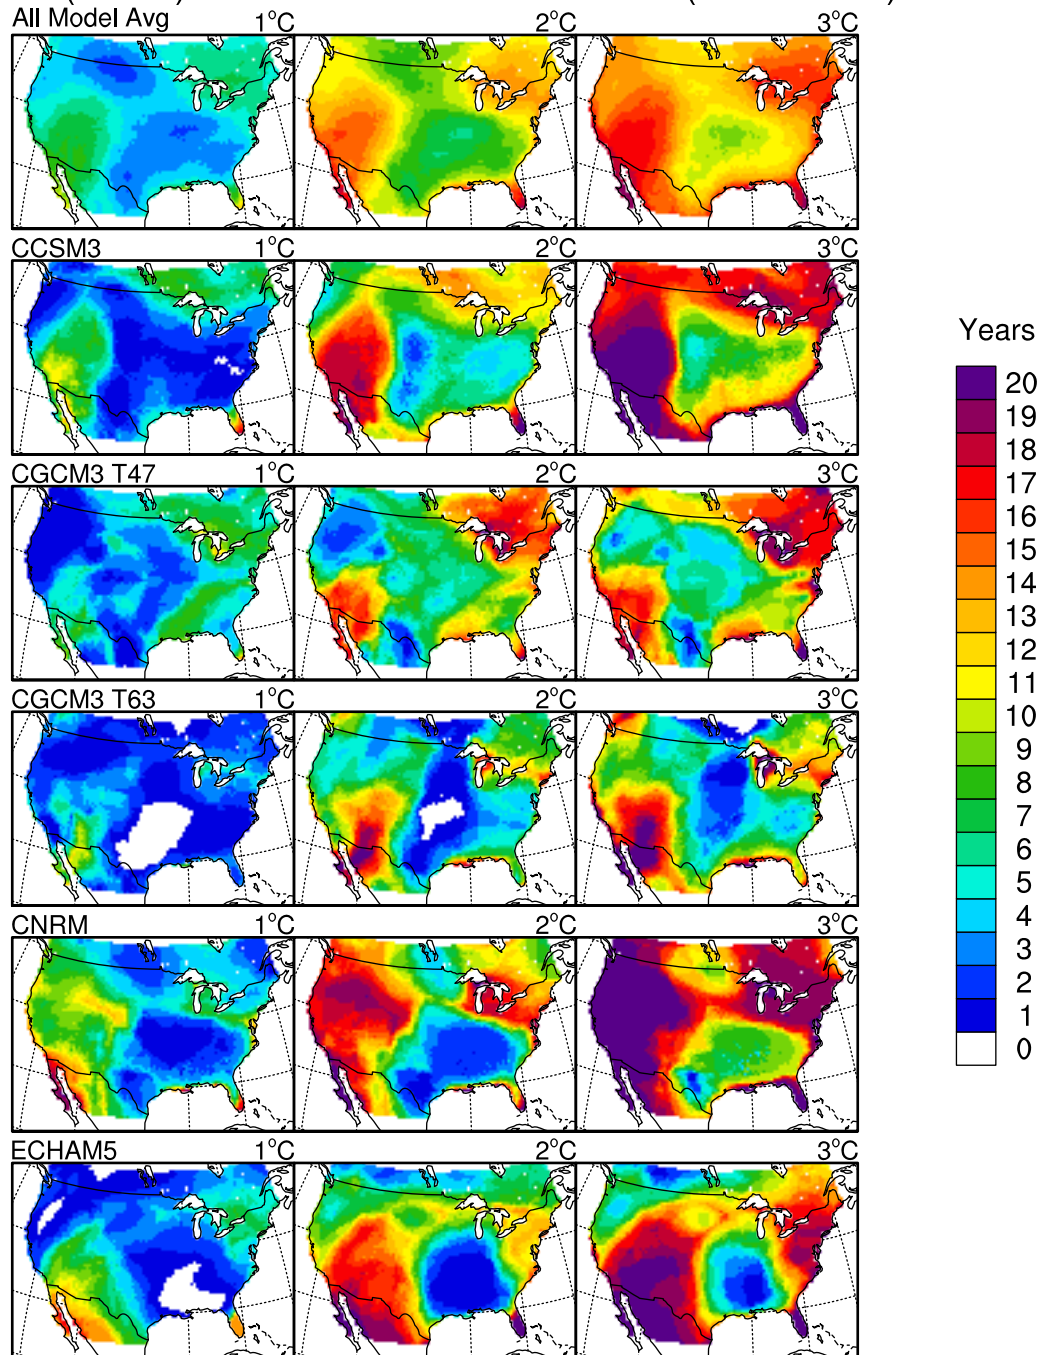
Avg Annual Cumulative Pr



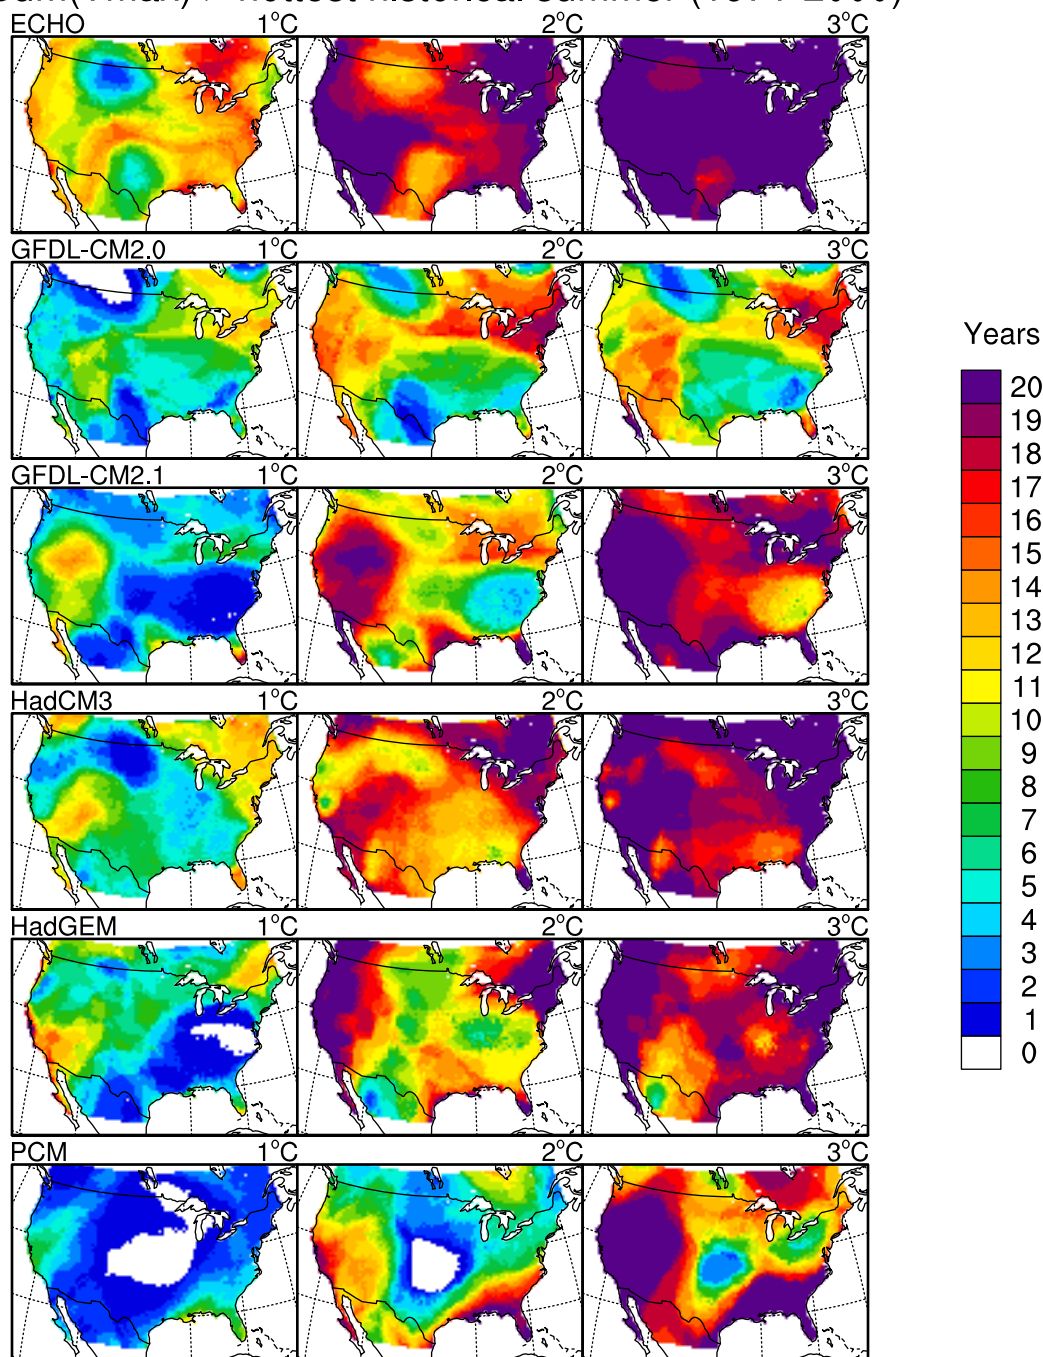
Avg Annual Cumulative Pr



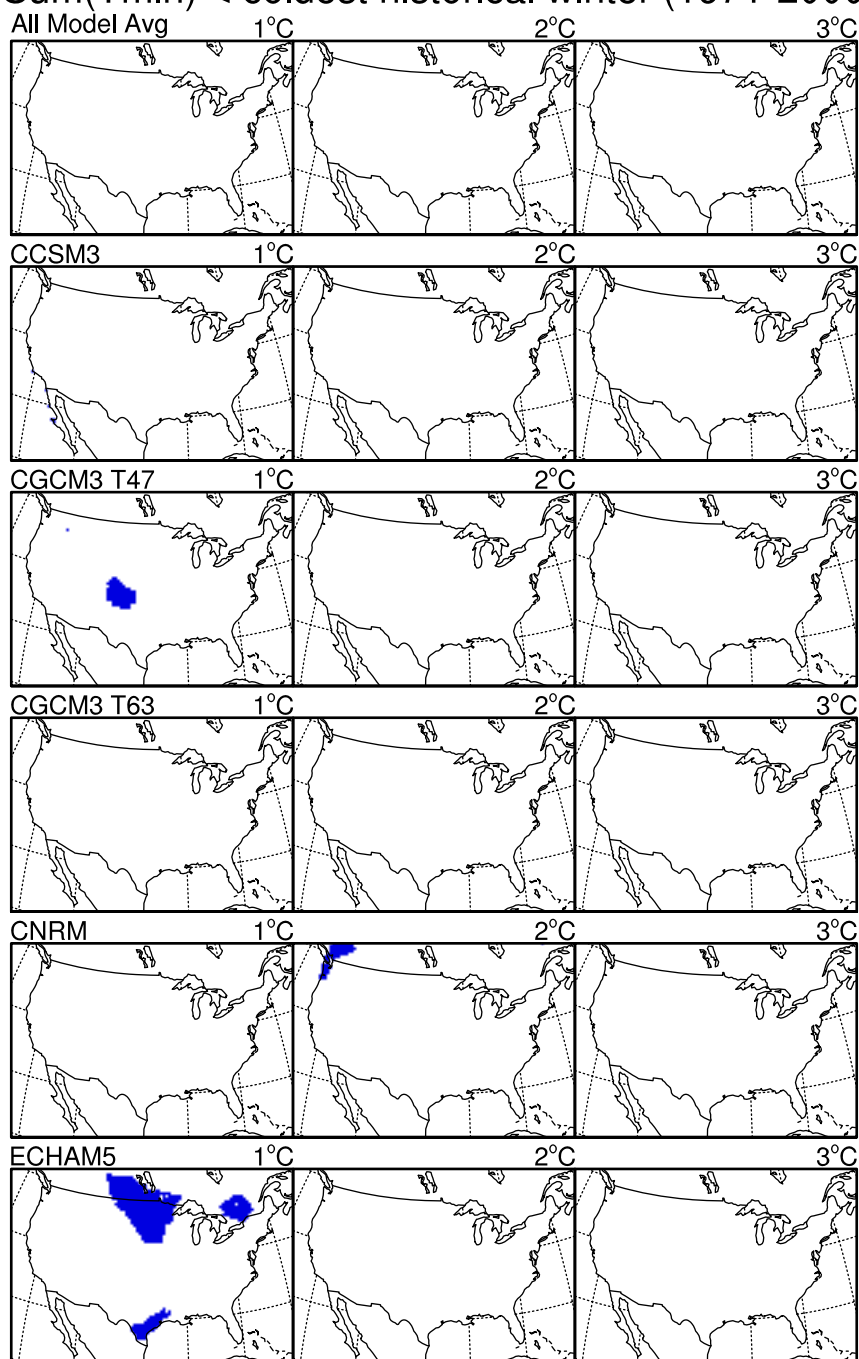
Hot summers: Number of yrs per 20yr period
Sum(Tmax) > hottest historical summer (1971-2000)




Hot summers: Number of yrs per 20yr period
Sum(Tmax) > hottest historical summer (1971-2000)



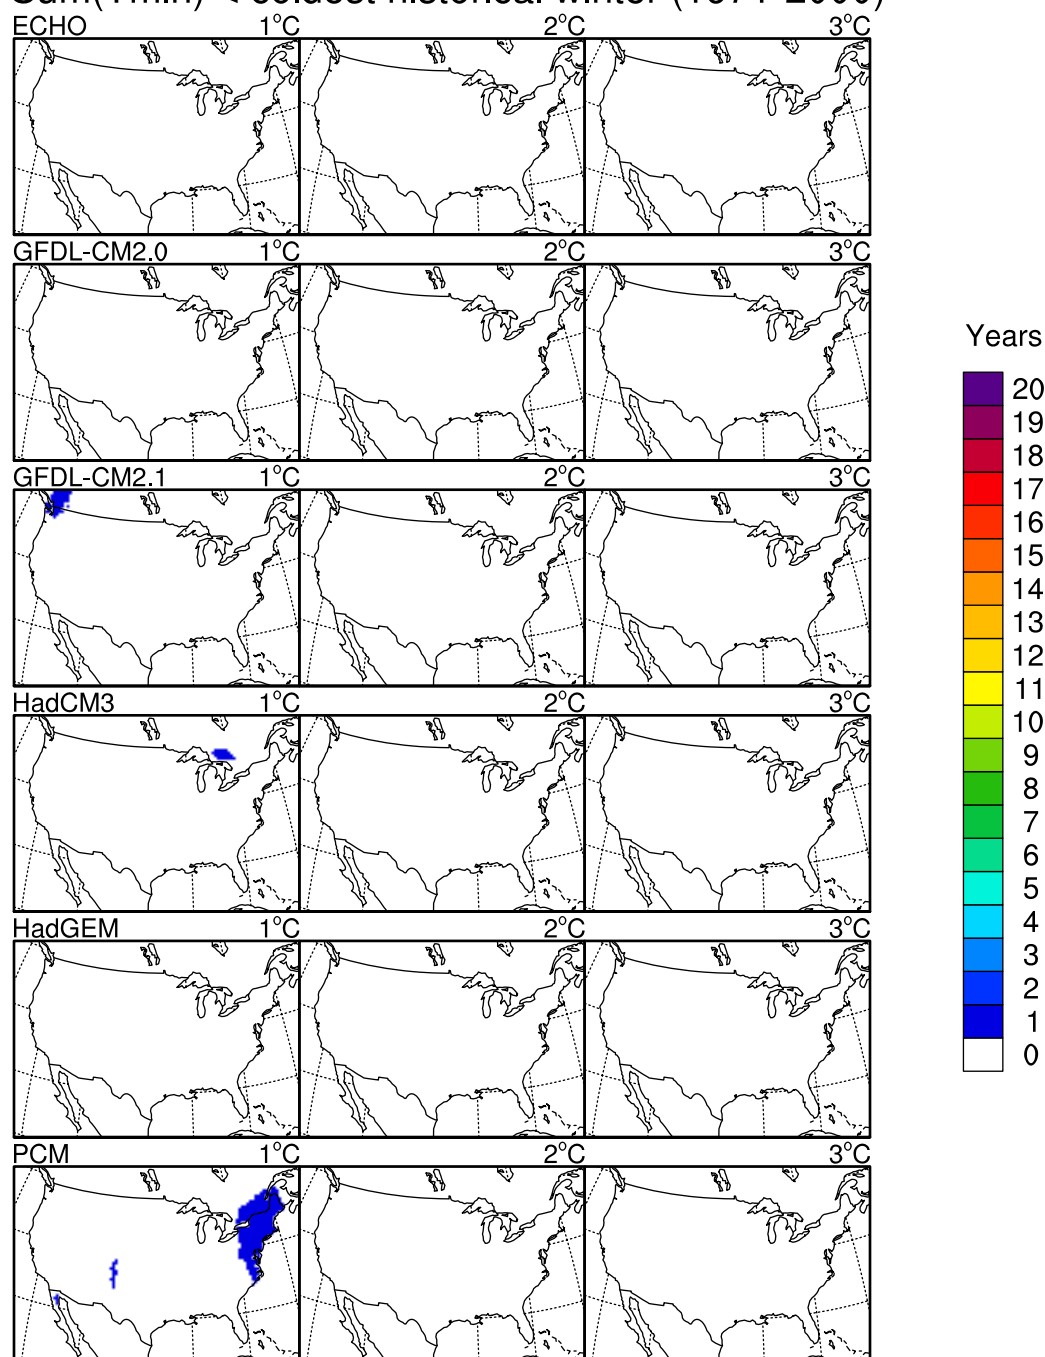
Cold winters: Number of yrs per 20yr period $\text{Sum}(\text{Tmin}) < \text{coldest historical winter (1971-2000)}$



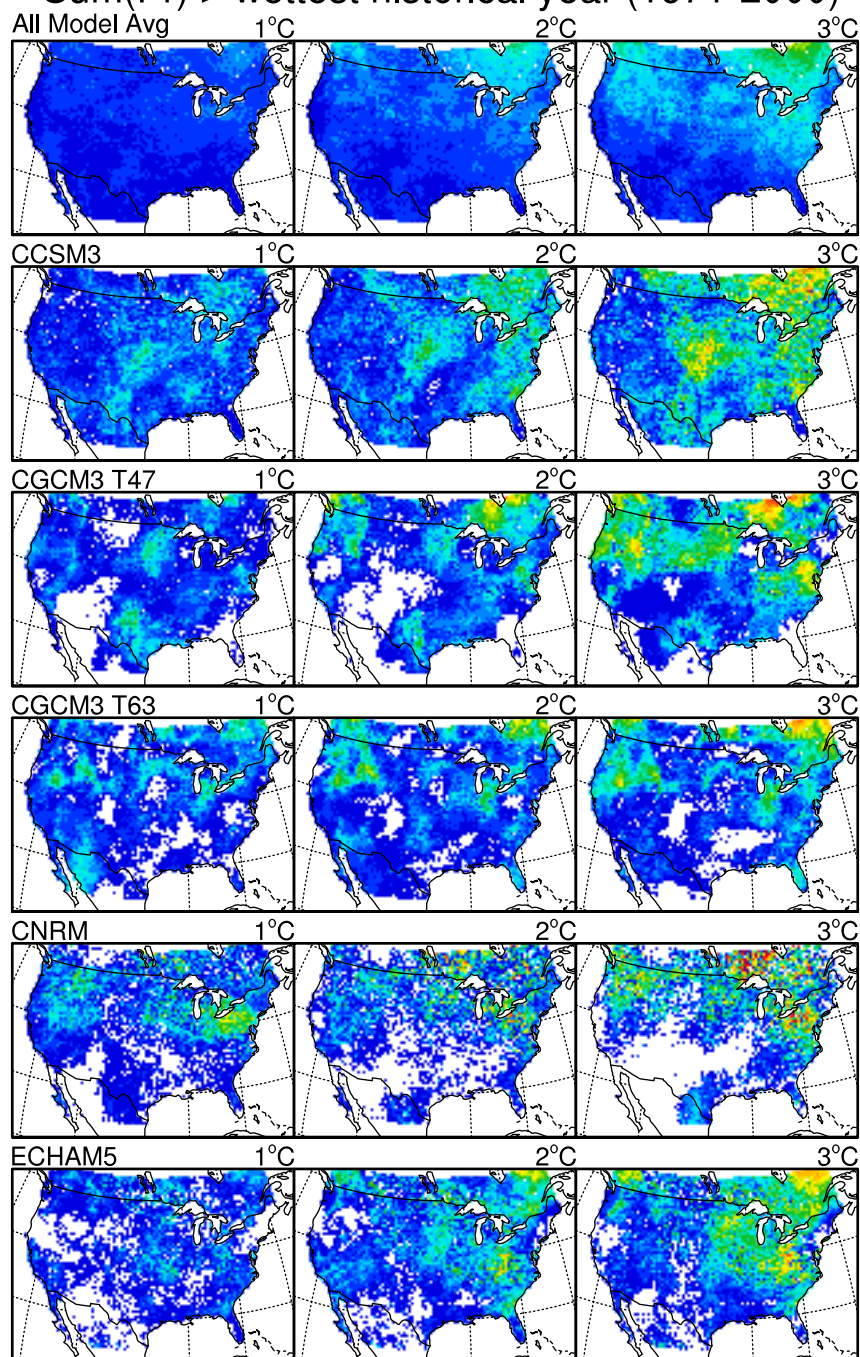
Years




Cold winters: Number of yrs per 20yr period
 $\text{Sum}(T_{\min}) < \text{coldest historical winter (1971-2000)}$



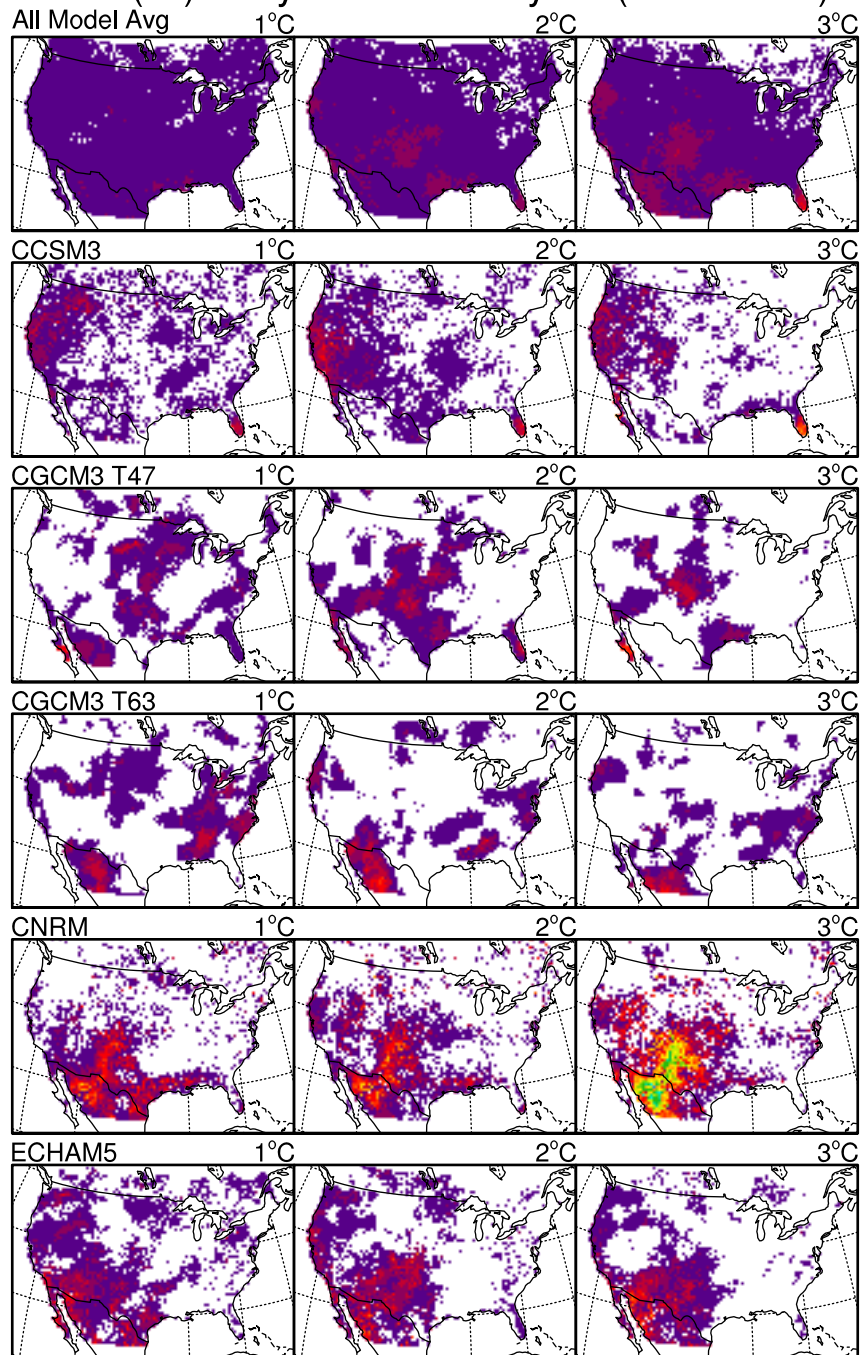
Wet years: Number of yrs per 20yr period
Sum(Pr) > wettest historical year (1971-2000)



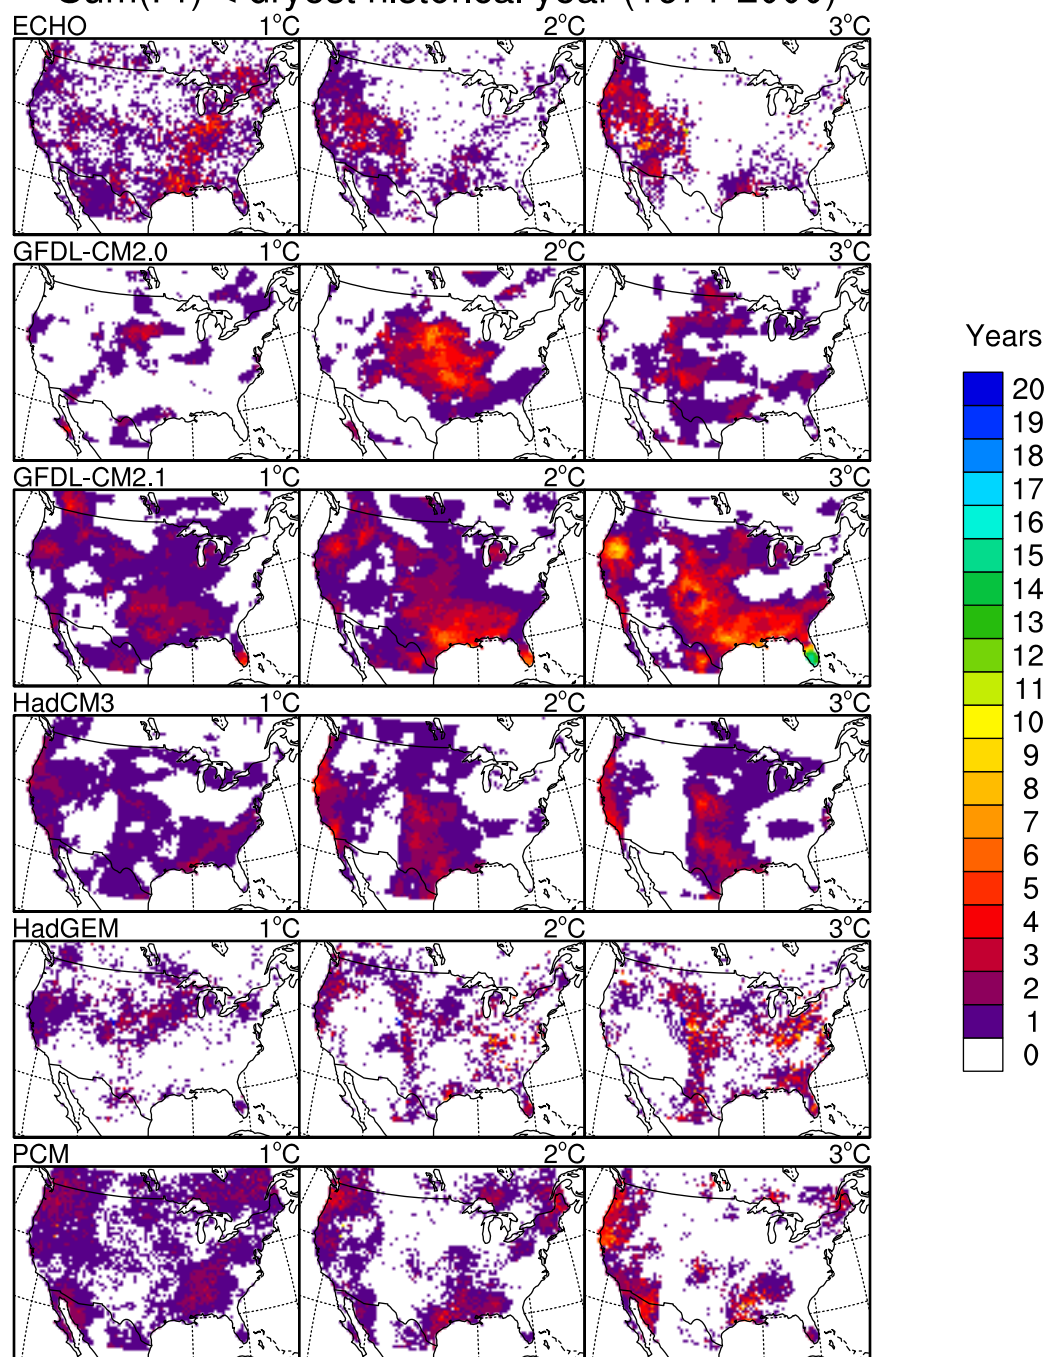
Wet years: Number of yrs per 20yr period
Sum(Pr) > wettest historical year (1971-2000)



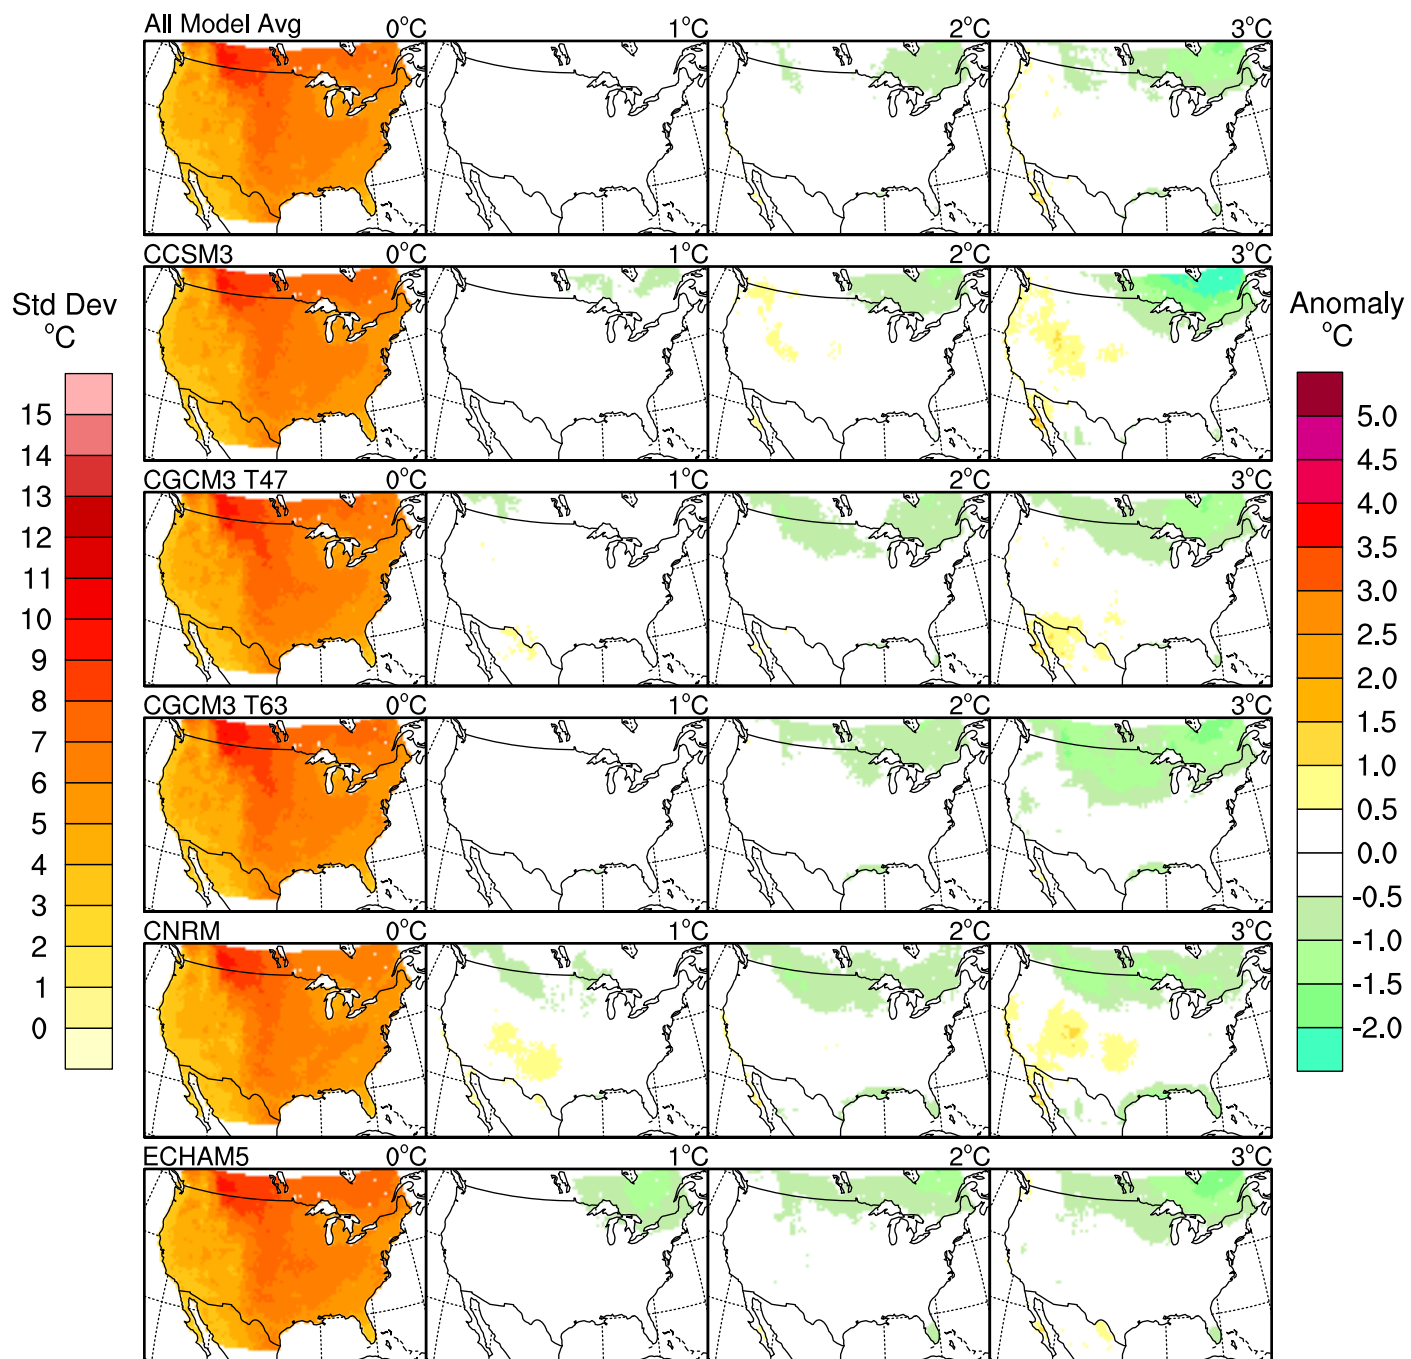
Dry years: Number of yrs per 20yr period
Sum(Pr) < dryest historical year (1971-2000)



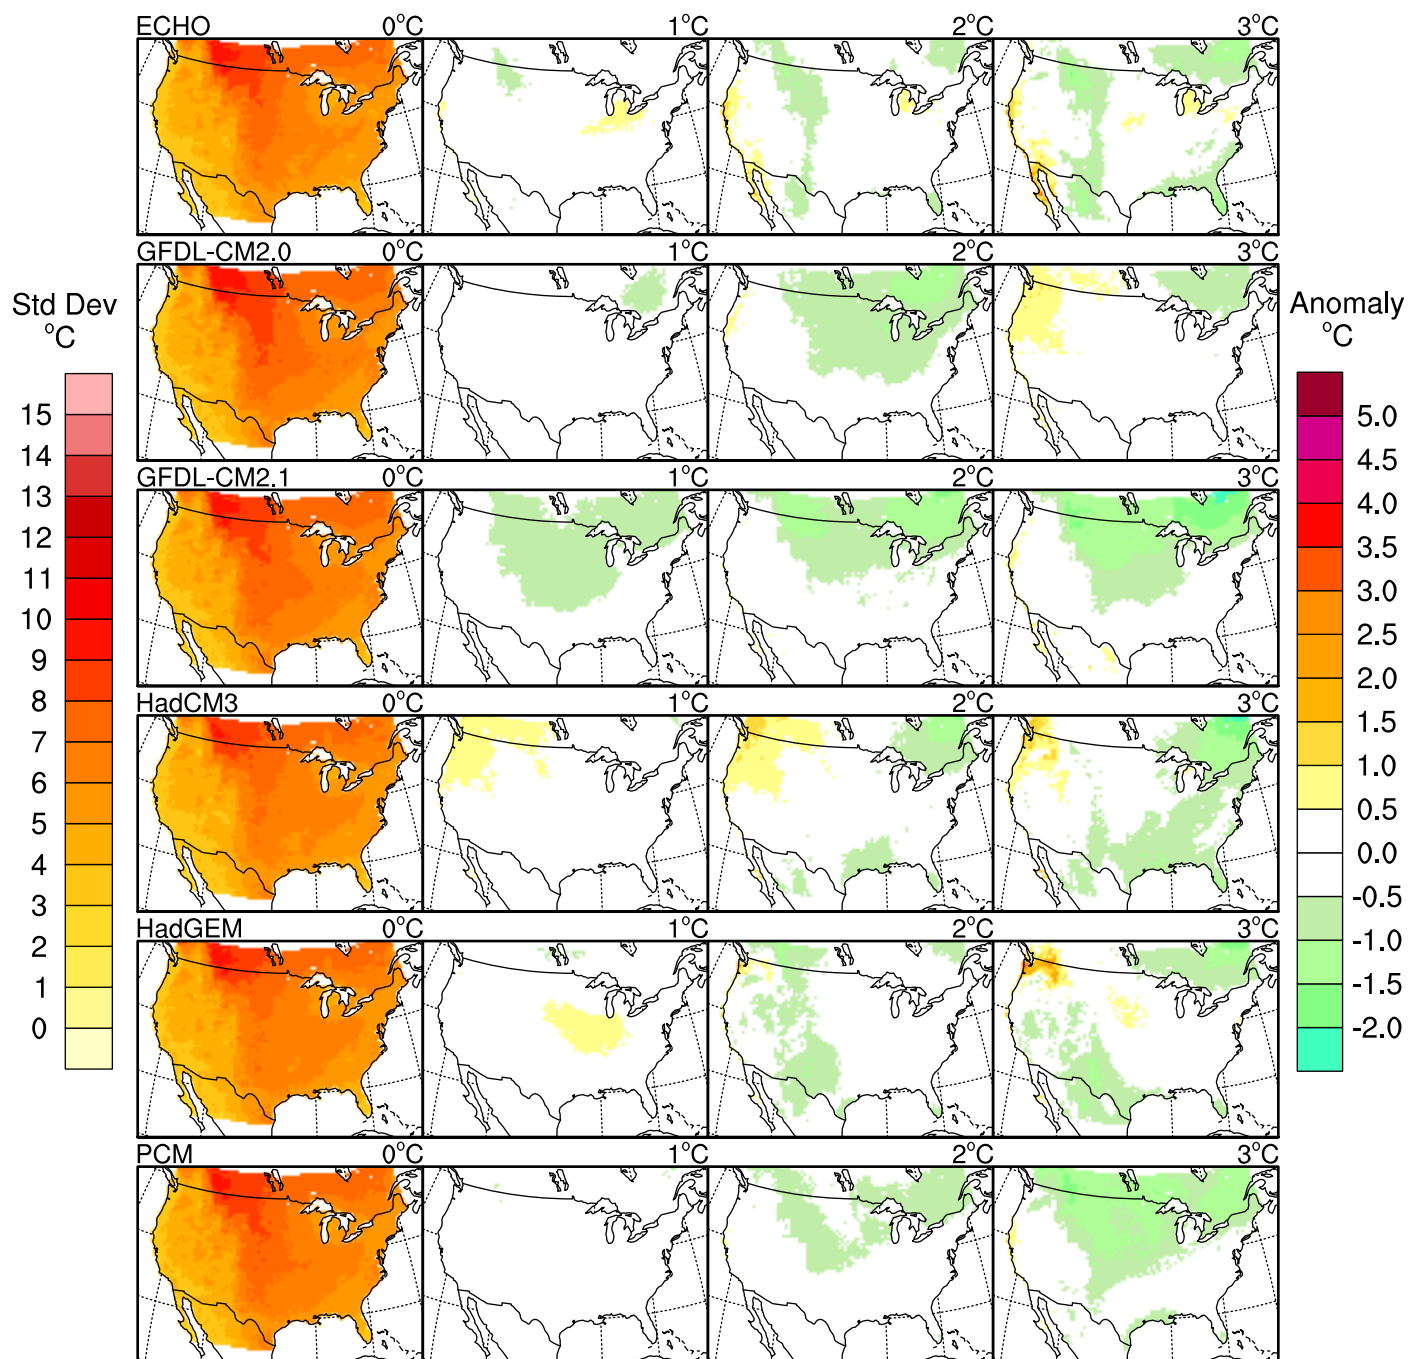
Dry years: Number of yrs per 20yr period
 $\text{Sum}(\text{Pr}) < \text{dryest historical year (1971-2000)}$



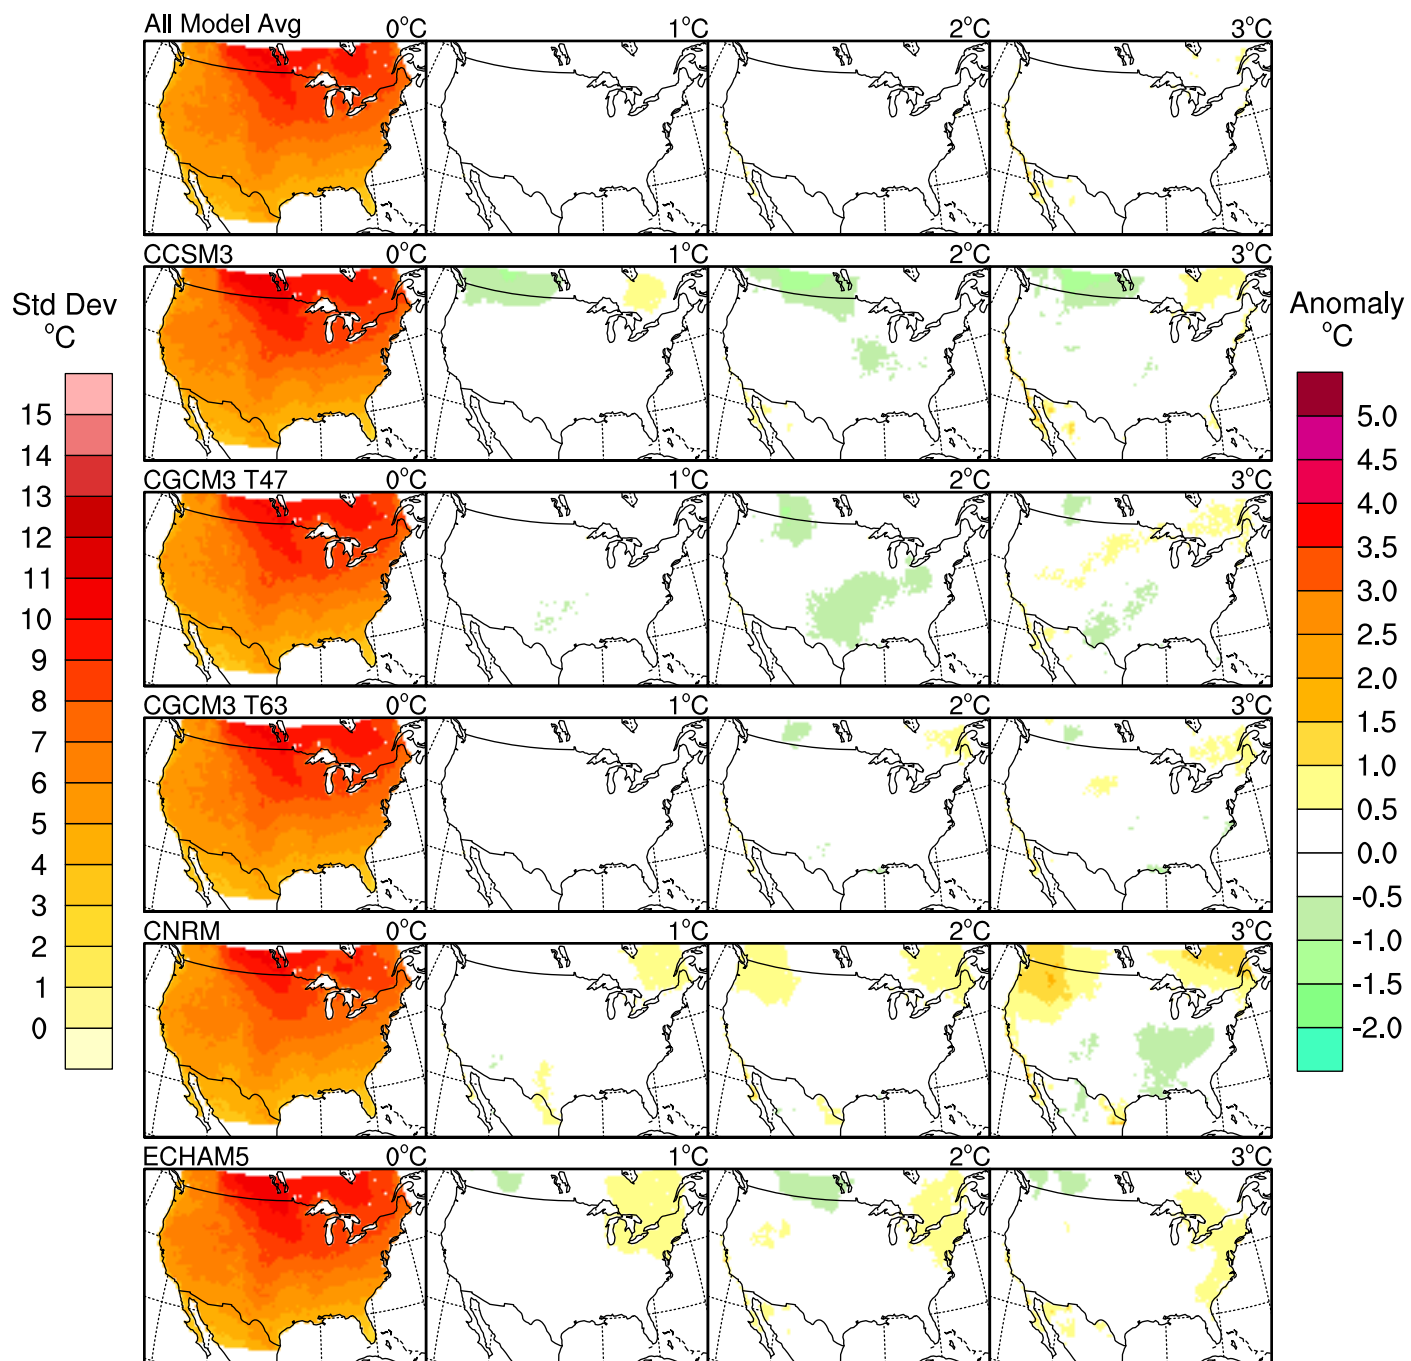
Winter (DJF) Tmax Standard Deviation




Winter (DJF) Tmax Standard Deviation




Spring (MAM) Tmax Standard Deviation



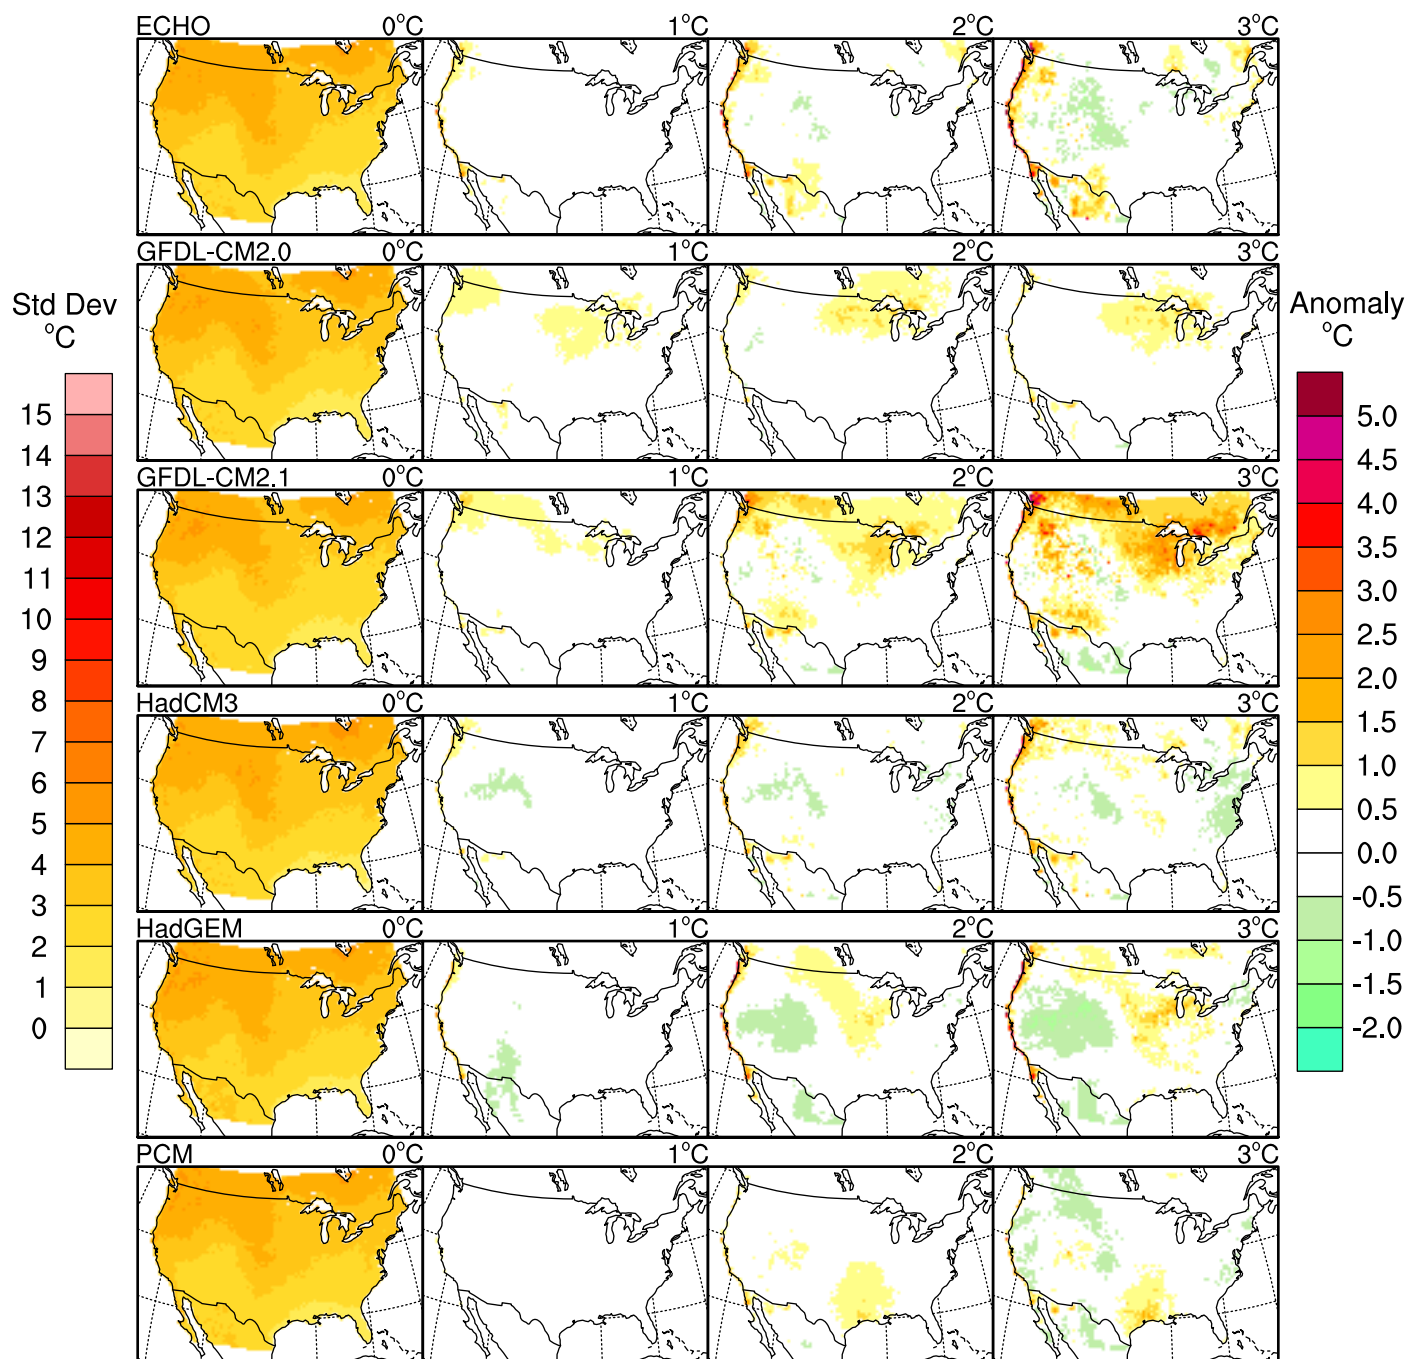
Spring (MAM) Tmax Standard Deviation



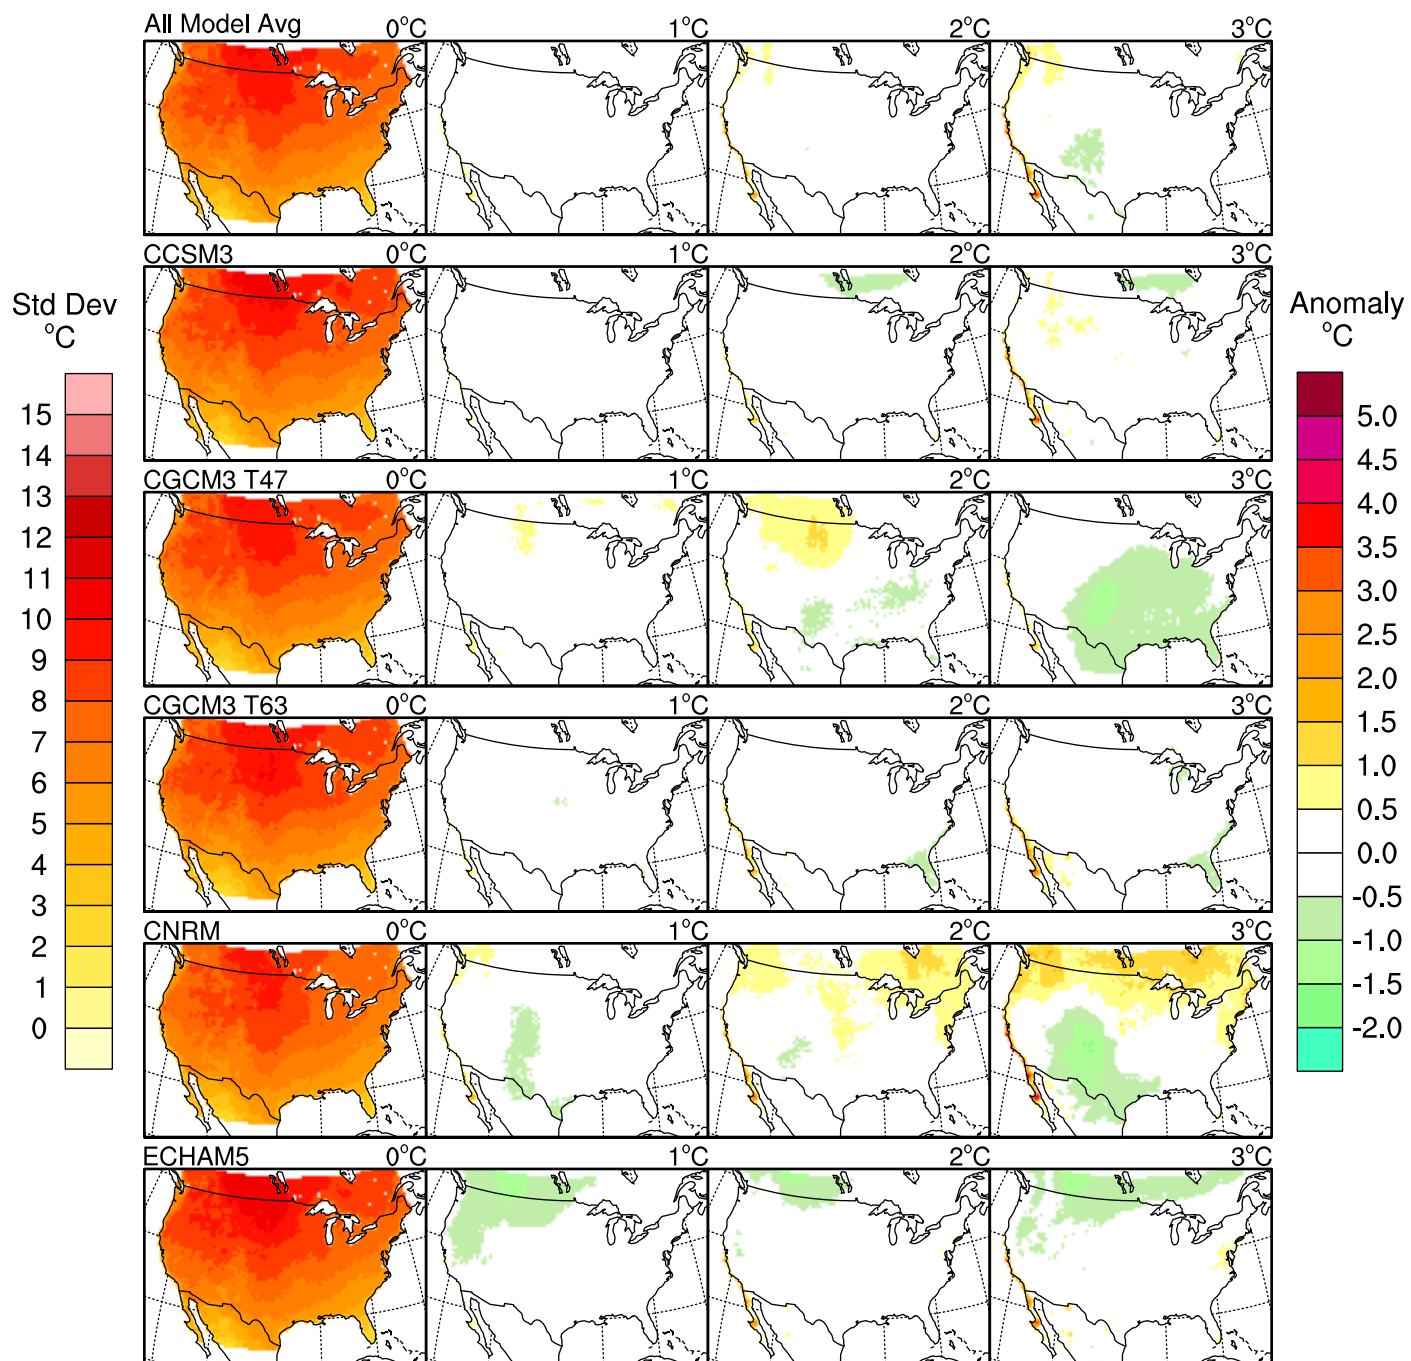
Summer (JJA) Tmax Standard Deviation



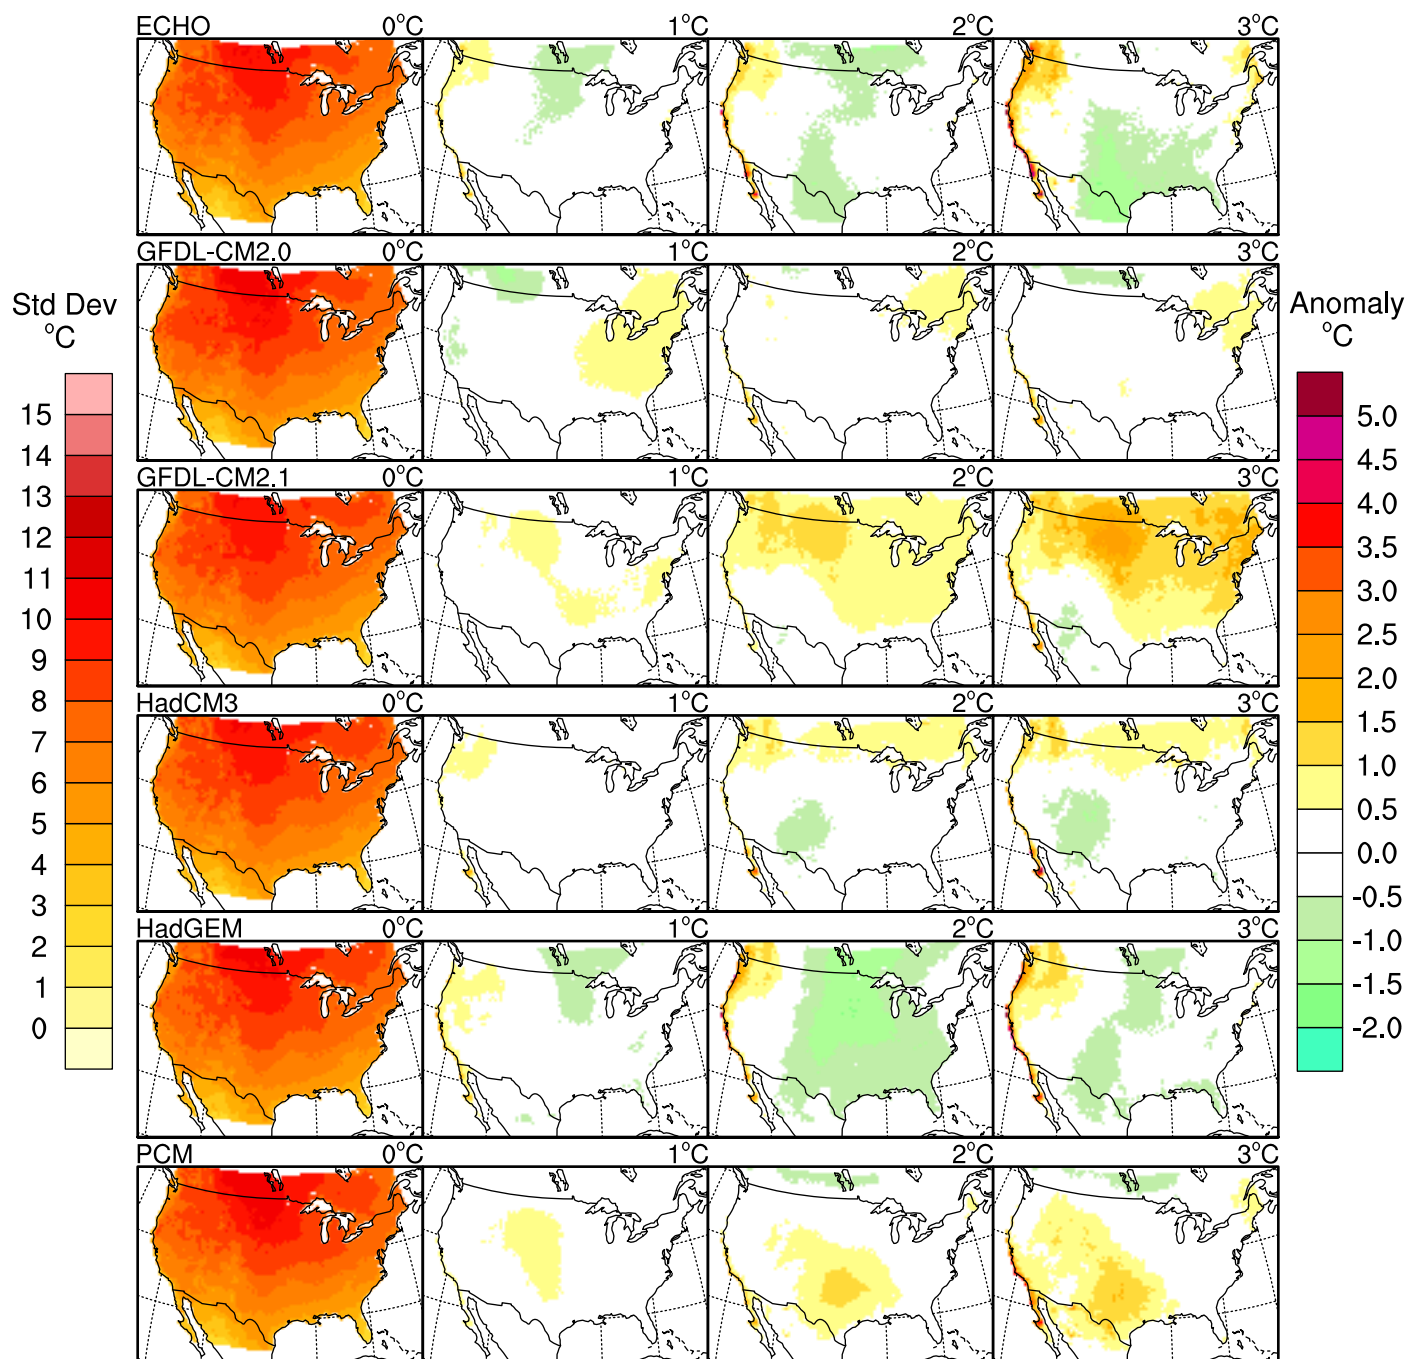
Summer (JJA) Tmax Standard Deviation



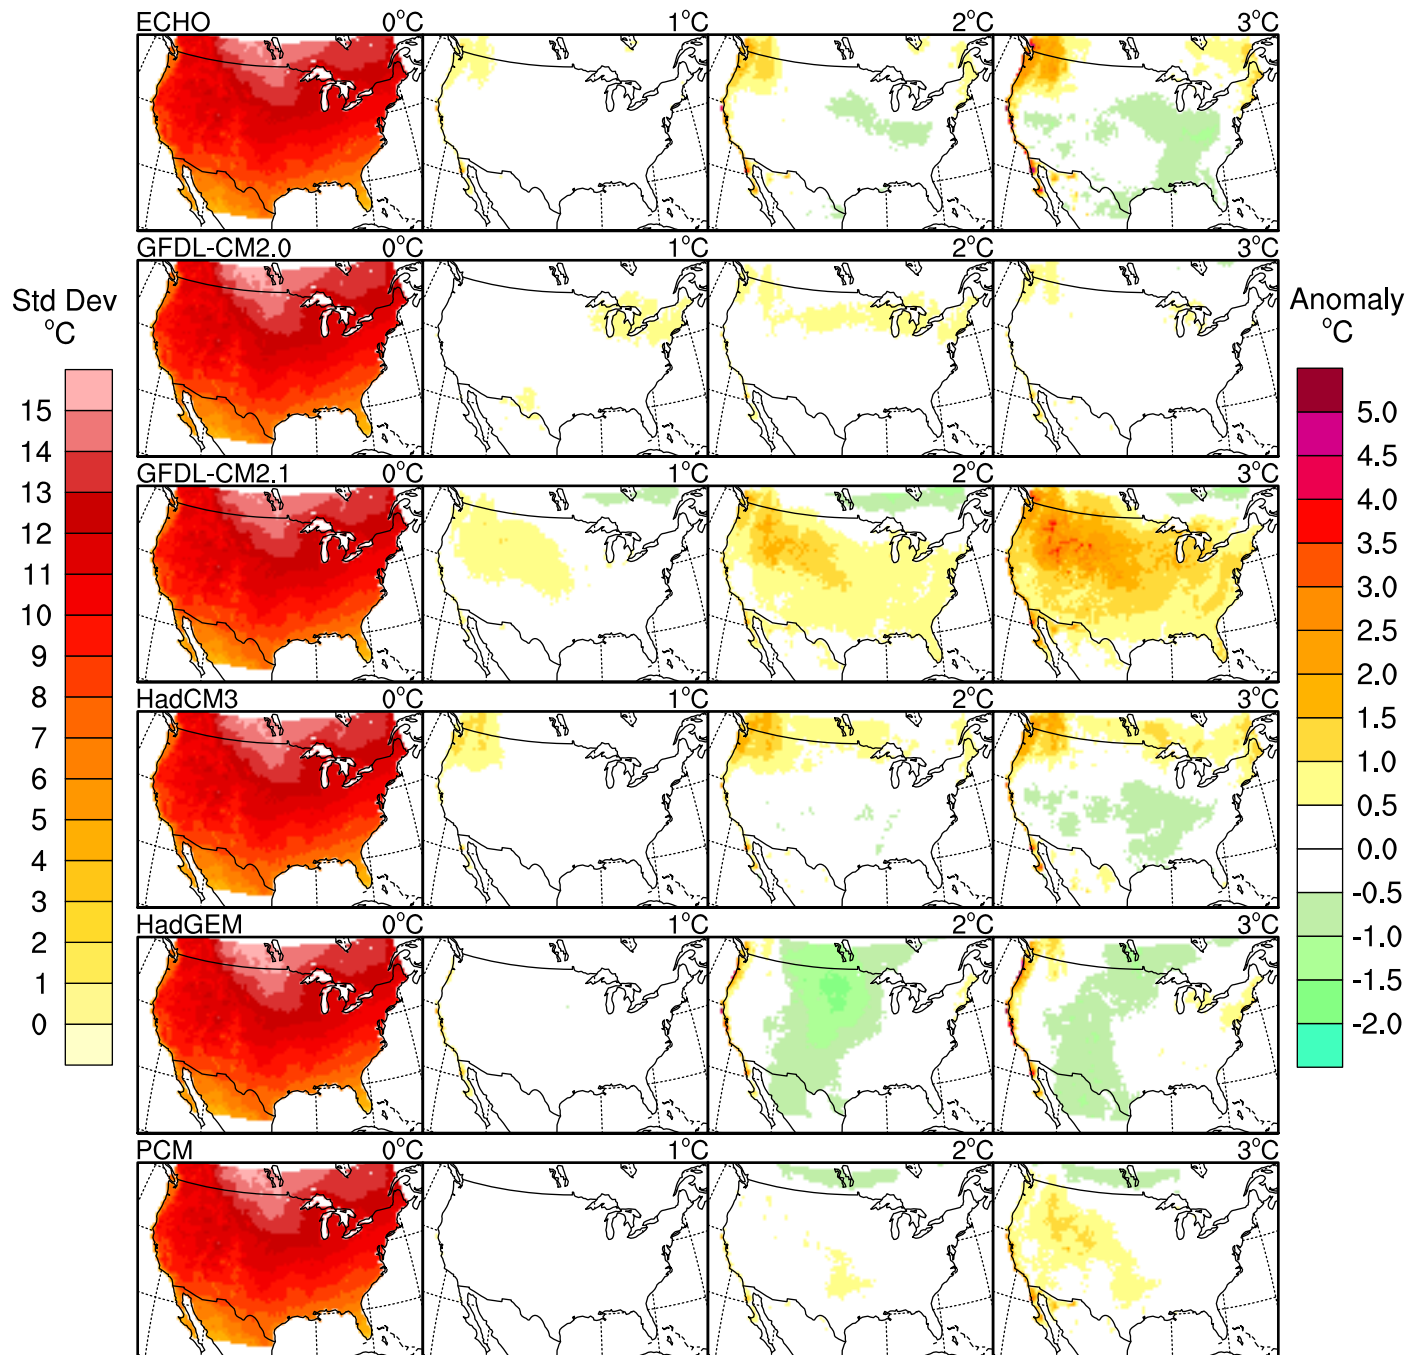
Fall (SON) Tmax Standard Deviation



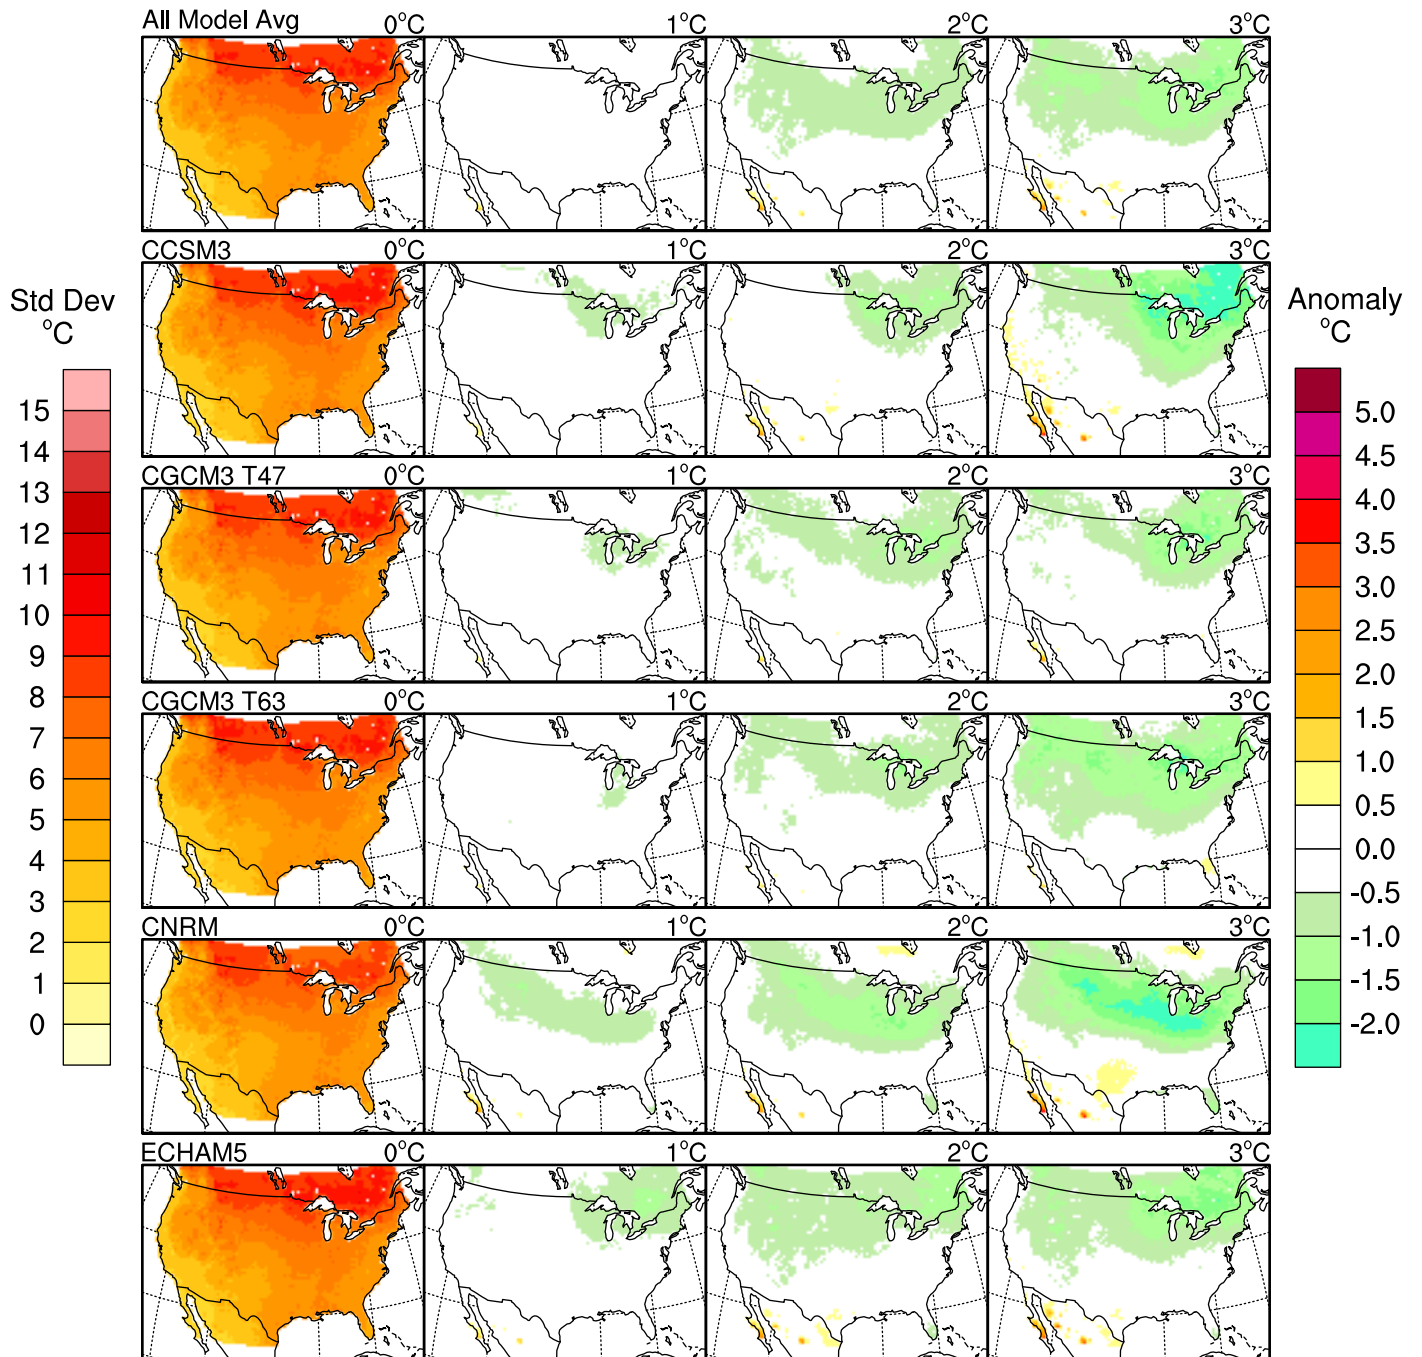
Fall (SON) Tmax Standard Deviation



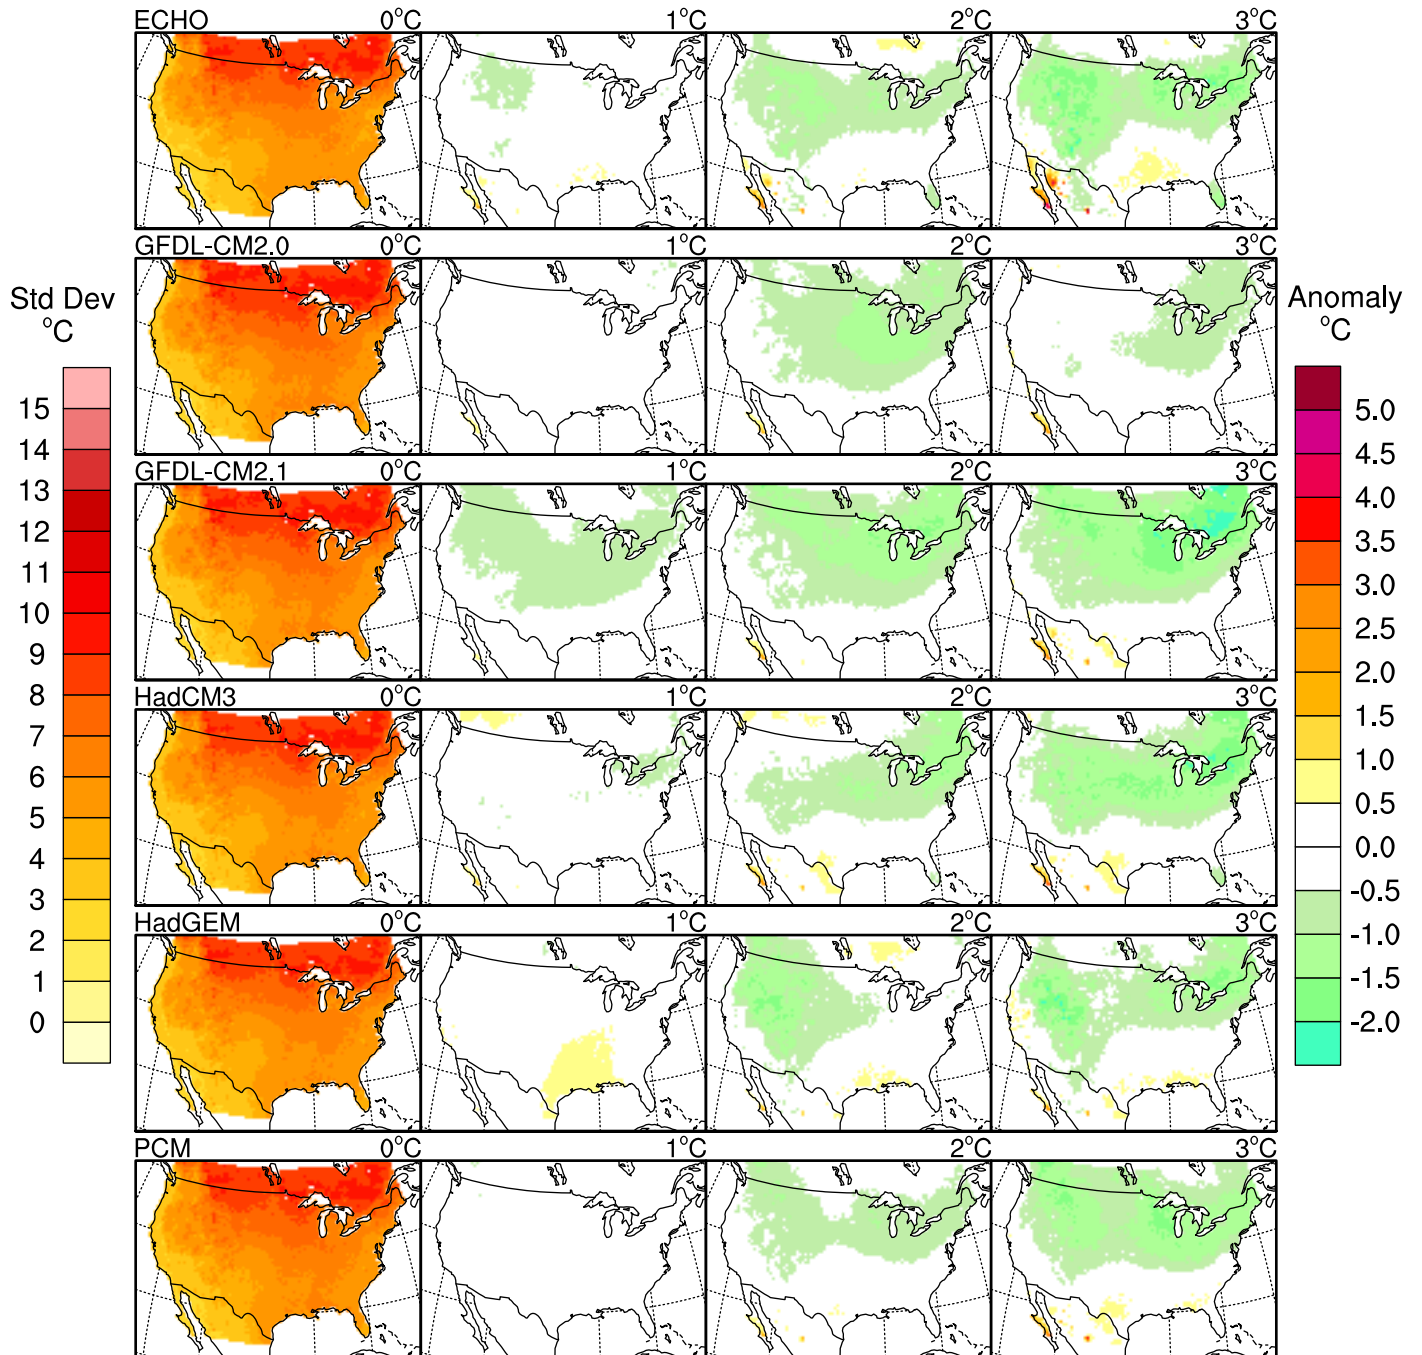
Annual Tmax Standard Deviation



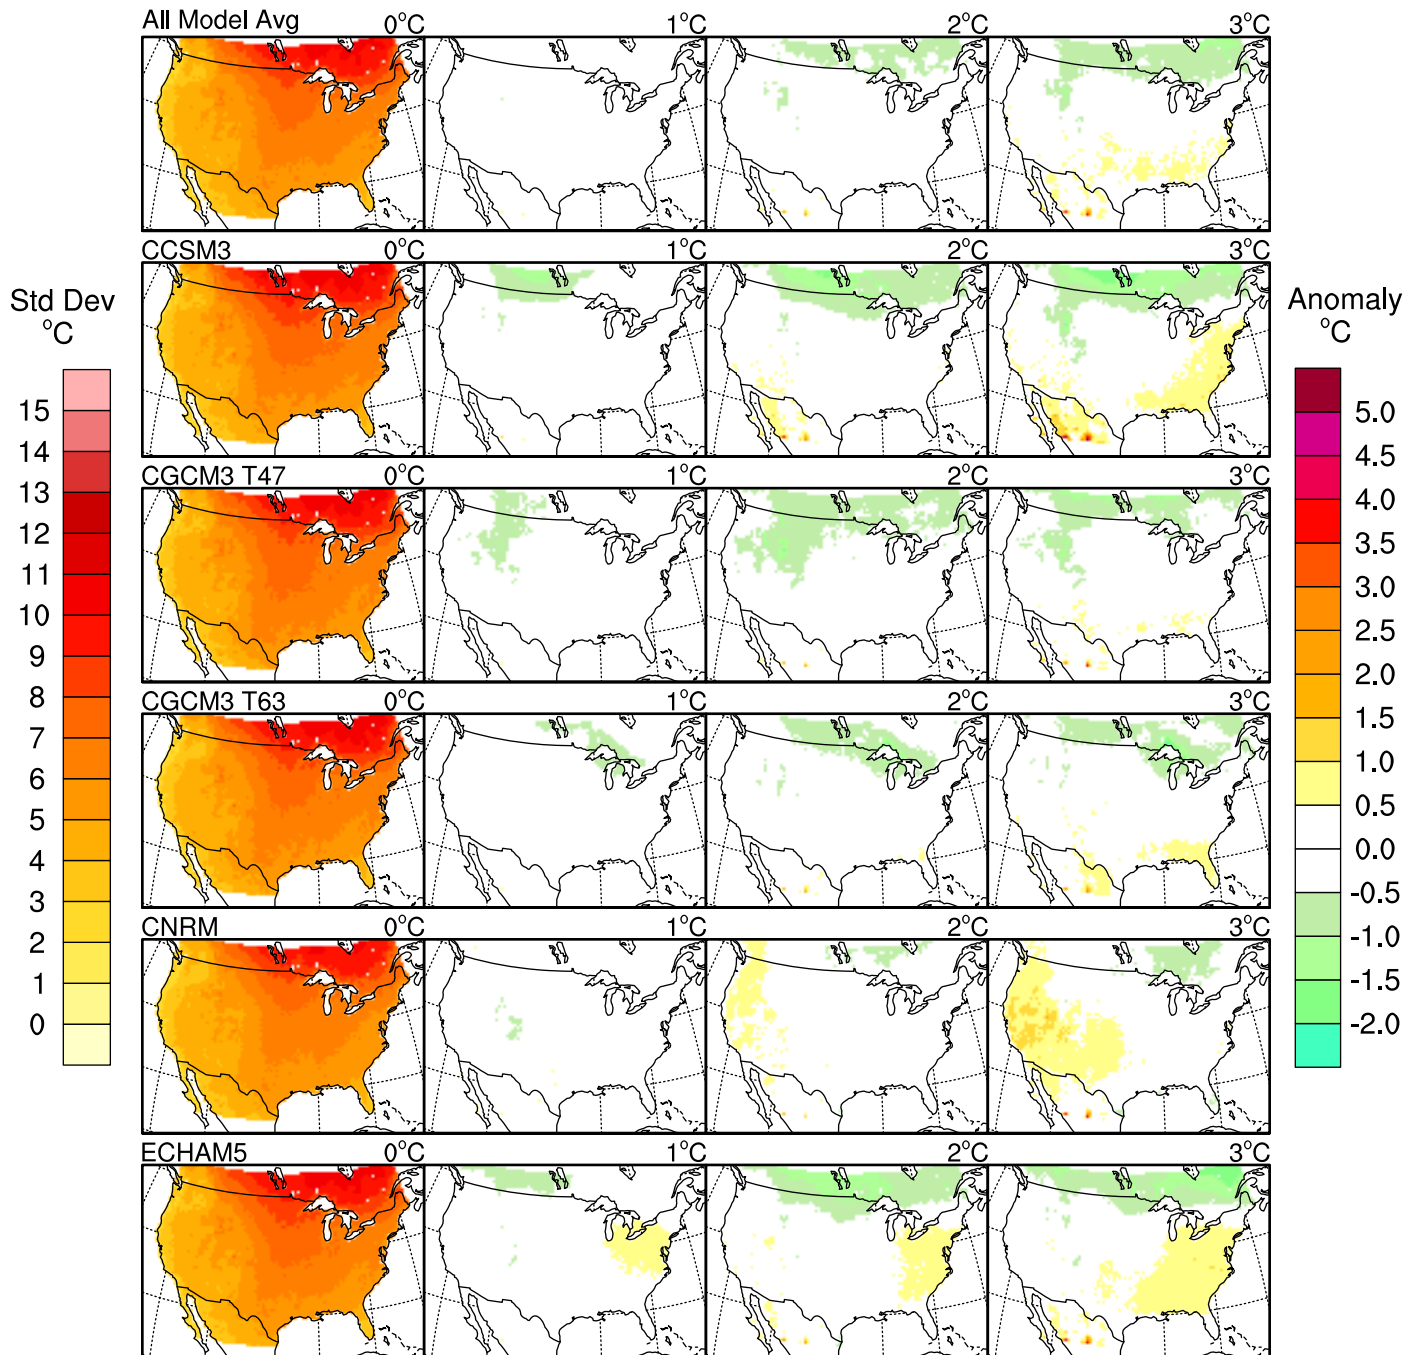
Winter (DJF) Tmin Standard Deviation



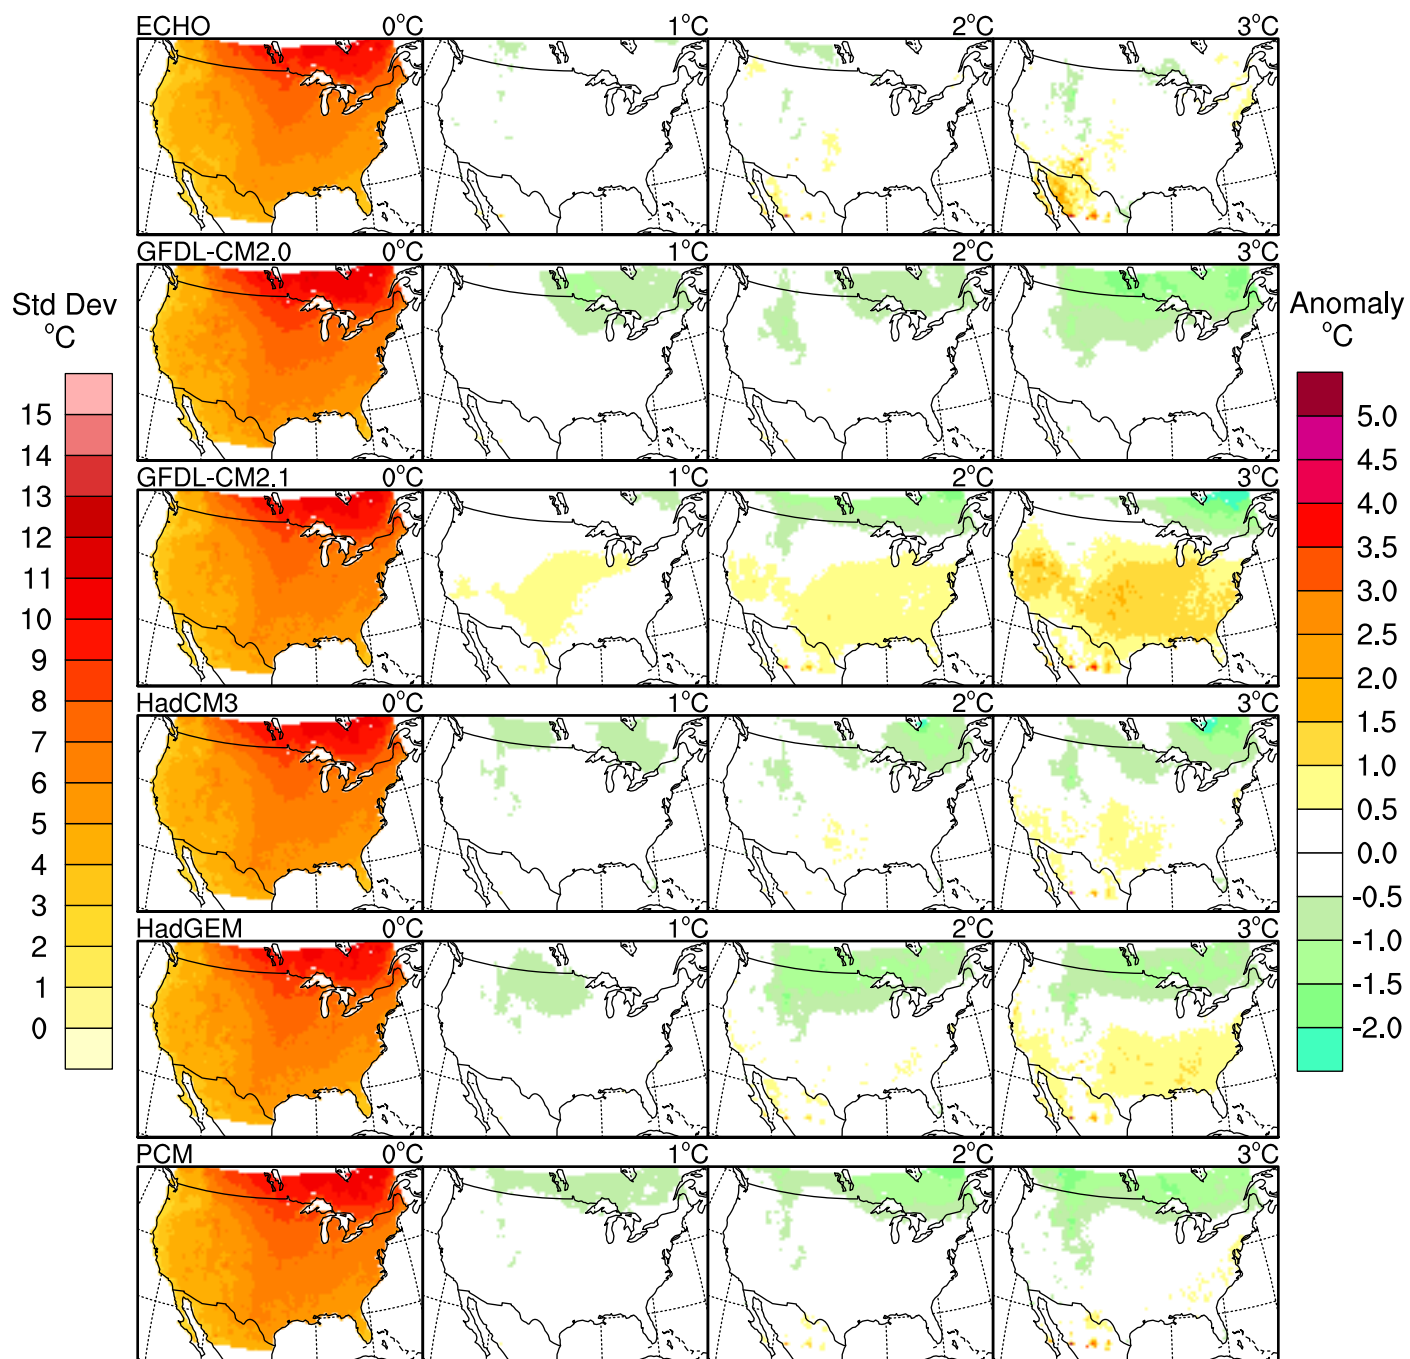
Winter (DJF) Tmin Standard Deviation




Spring (MAM) Tmin Standard Deviation



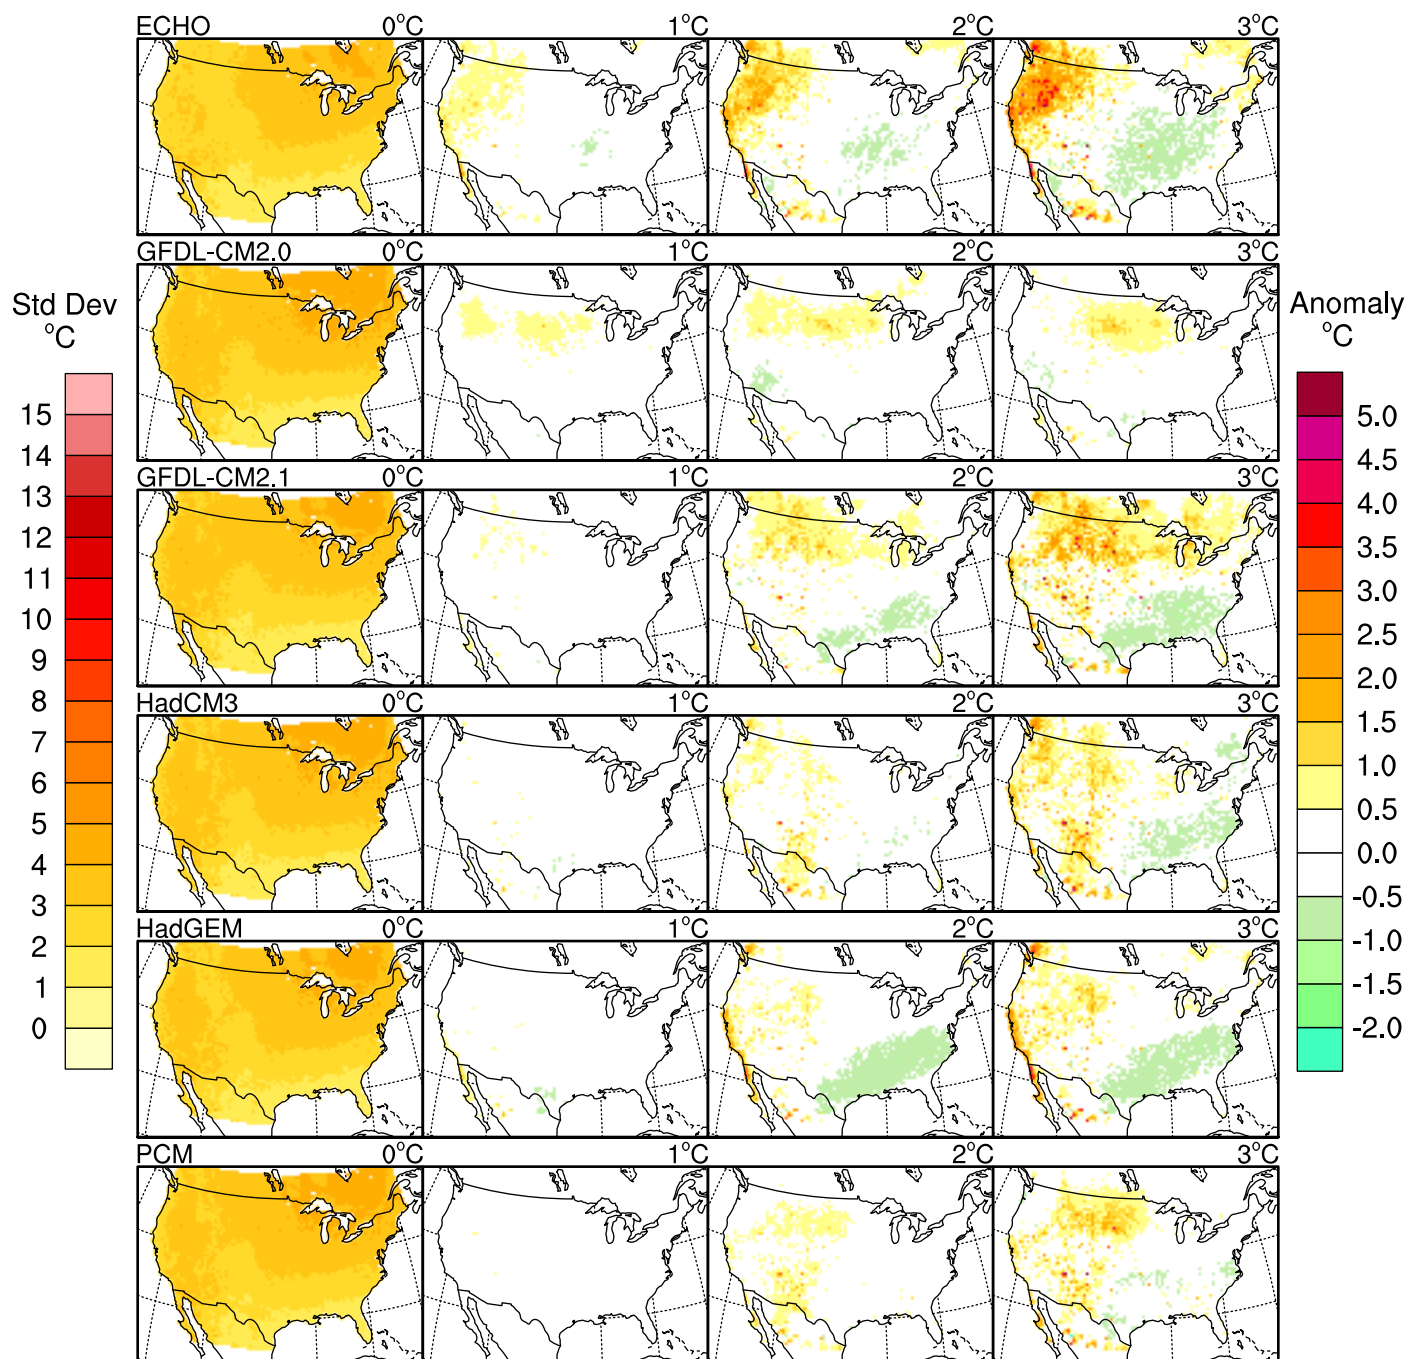
Spring (MAM) Tmin Standard Deviation



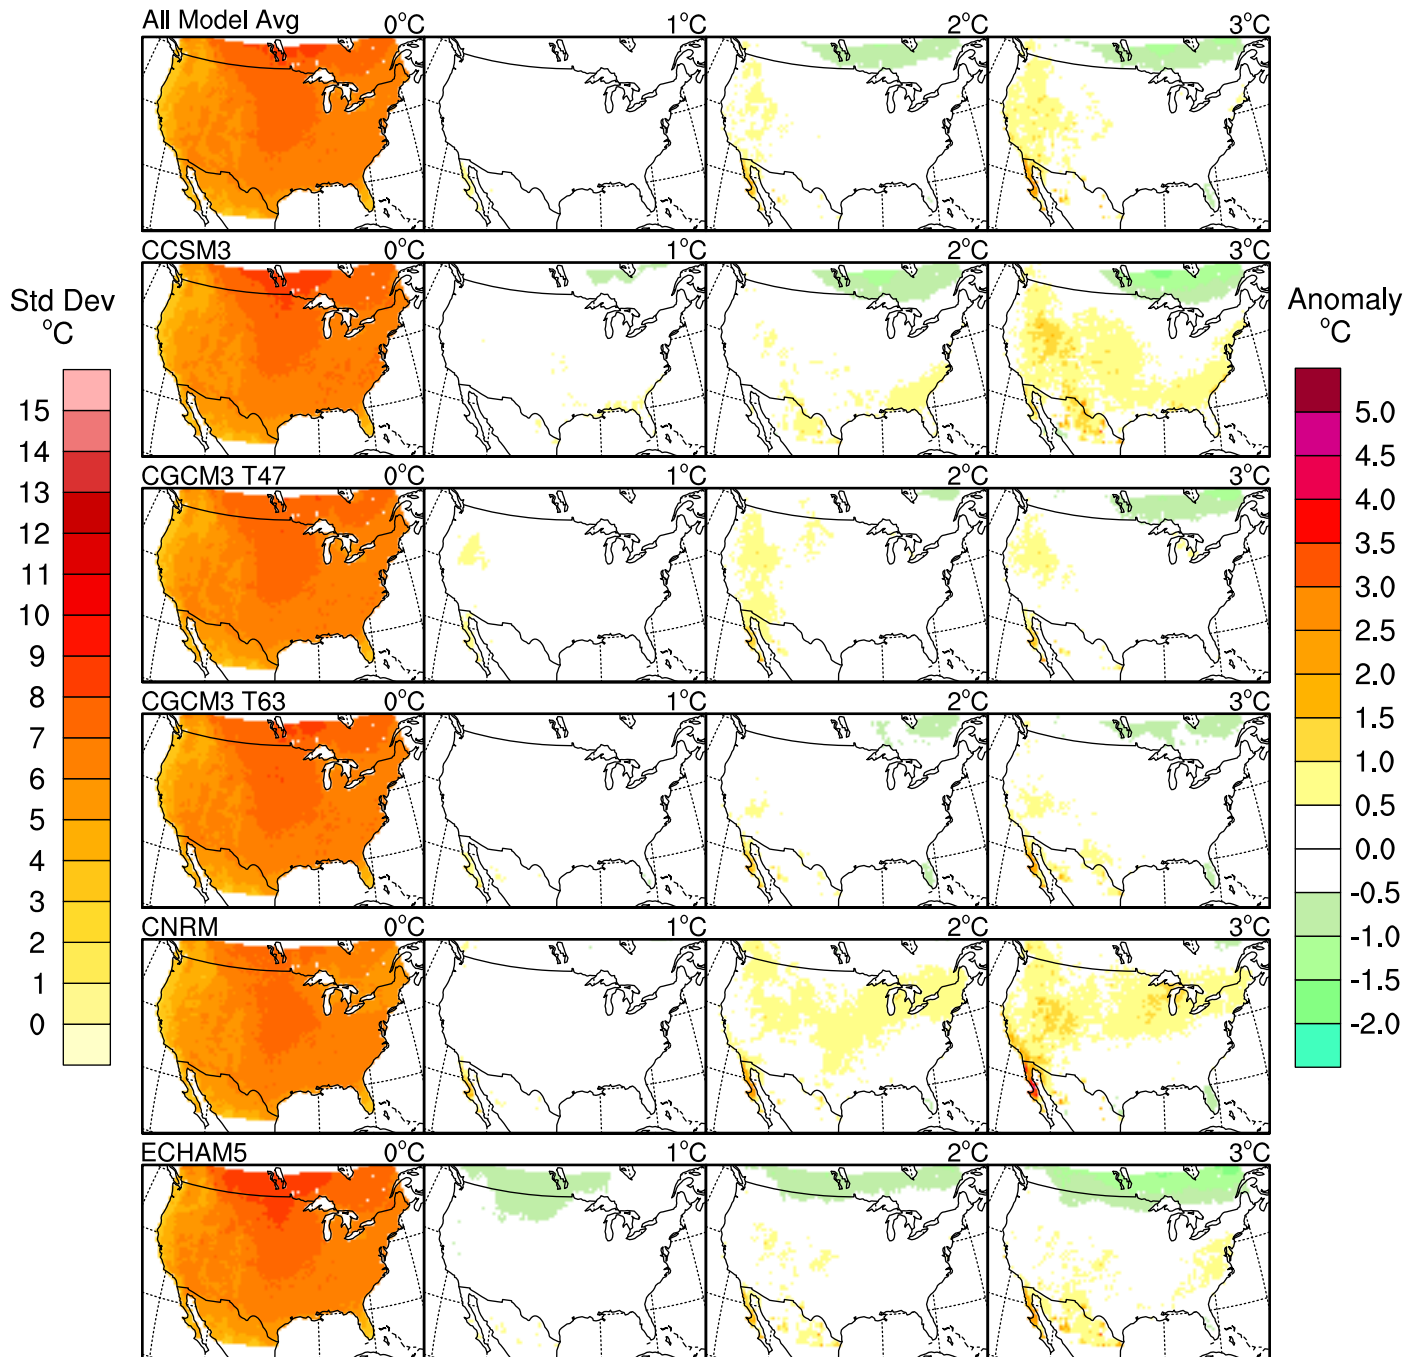
Summer (JJA) Tmin Standard Deviation



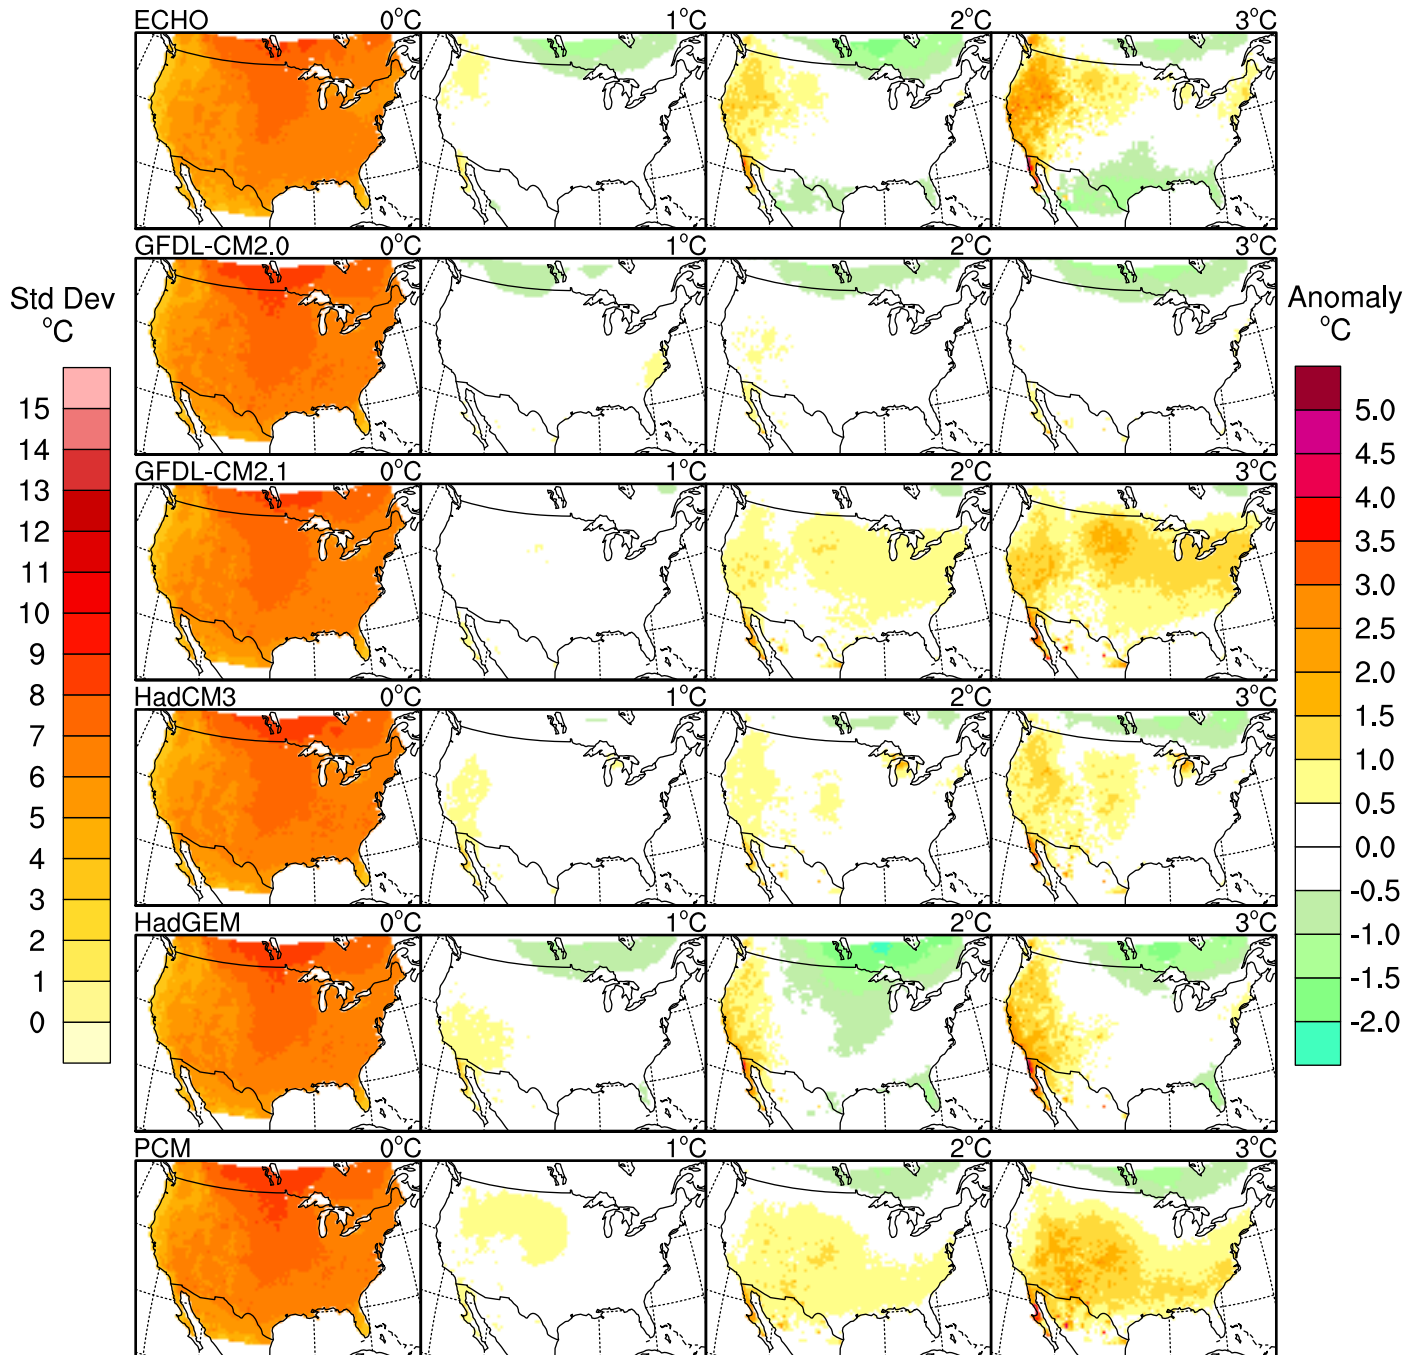
Summer (JJA) T_{min} Standard Deviation



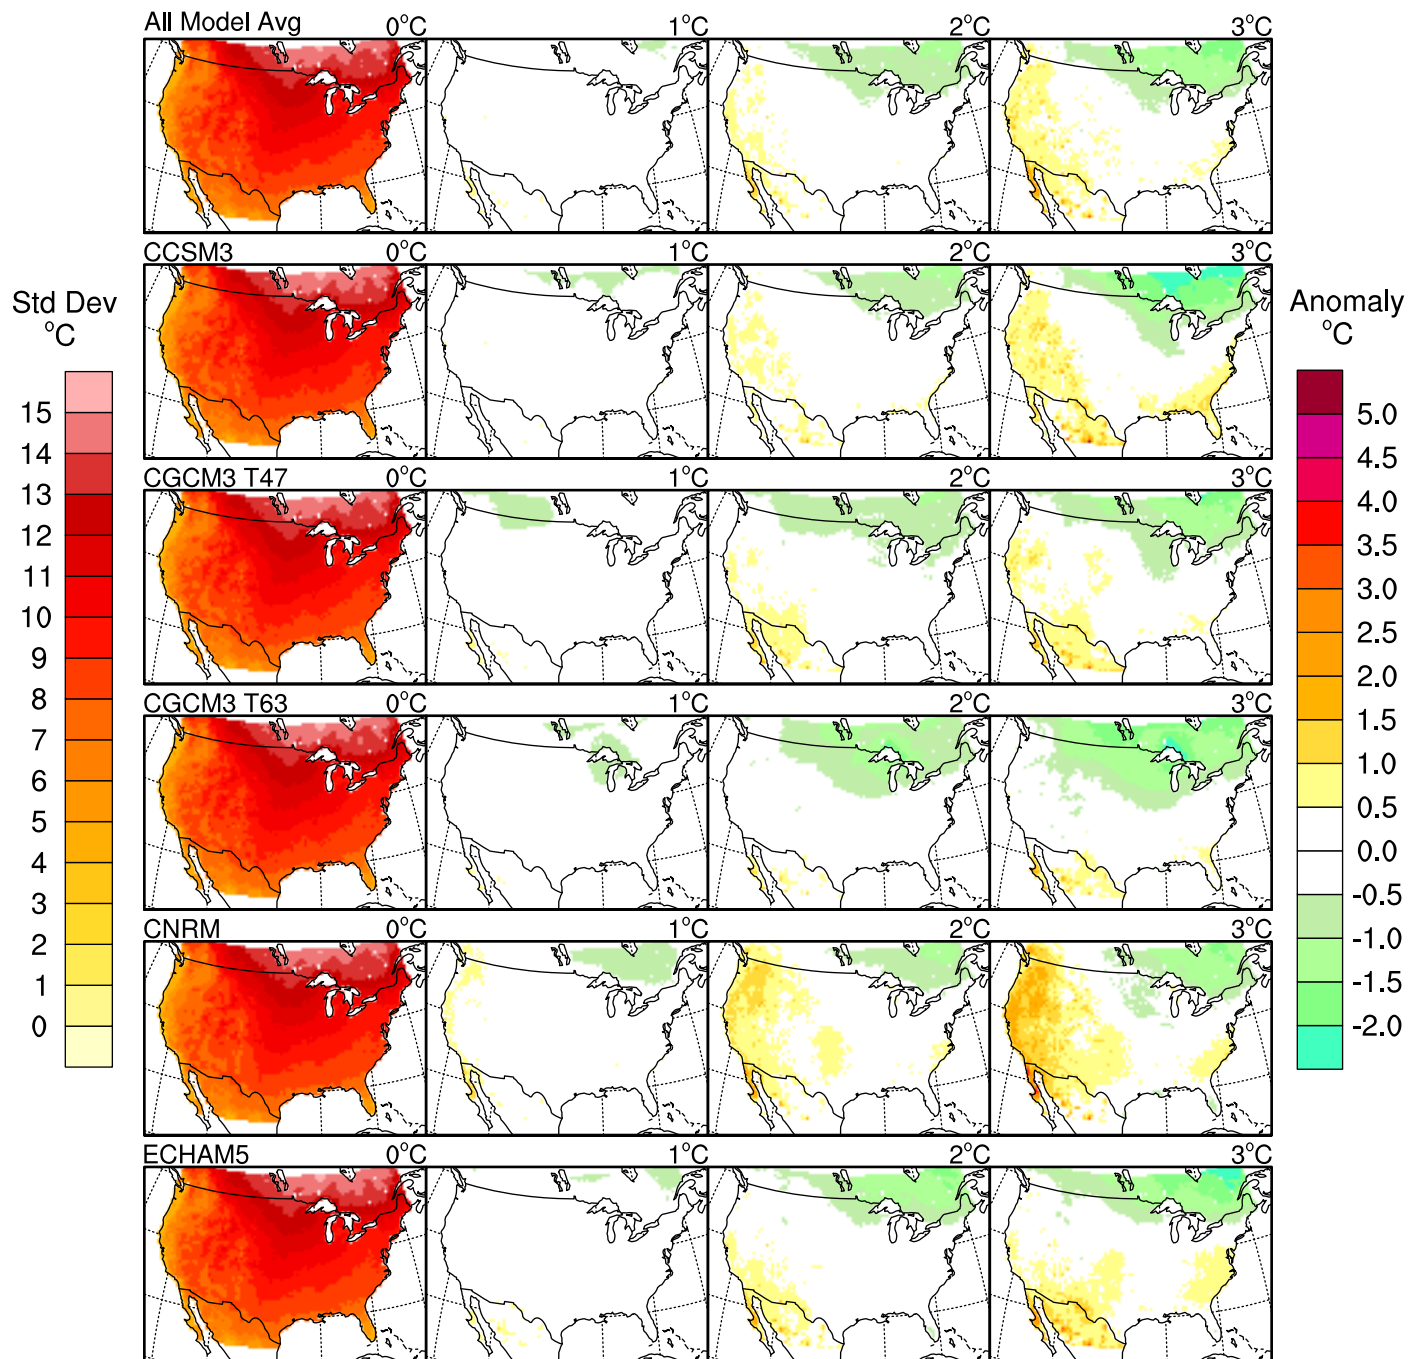
Fall (SON) Tmin Standard Deviation



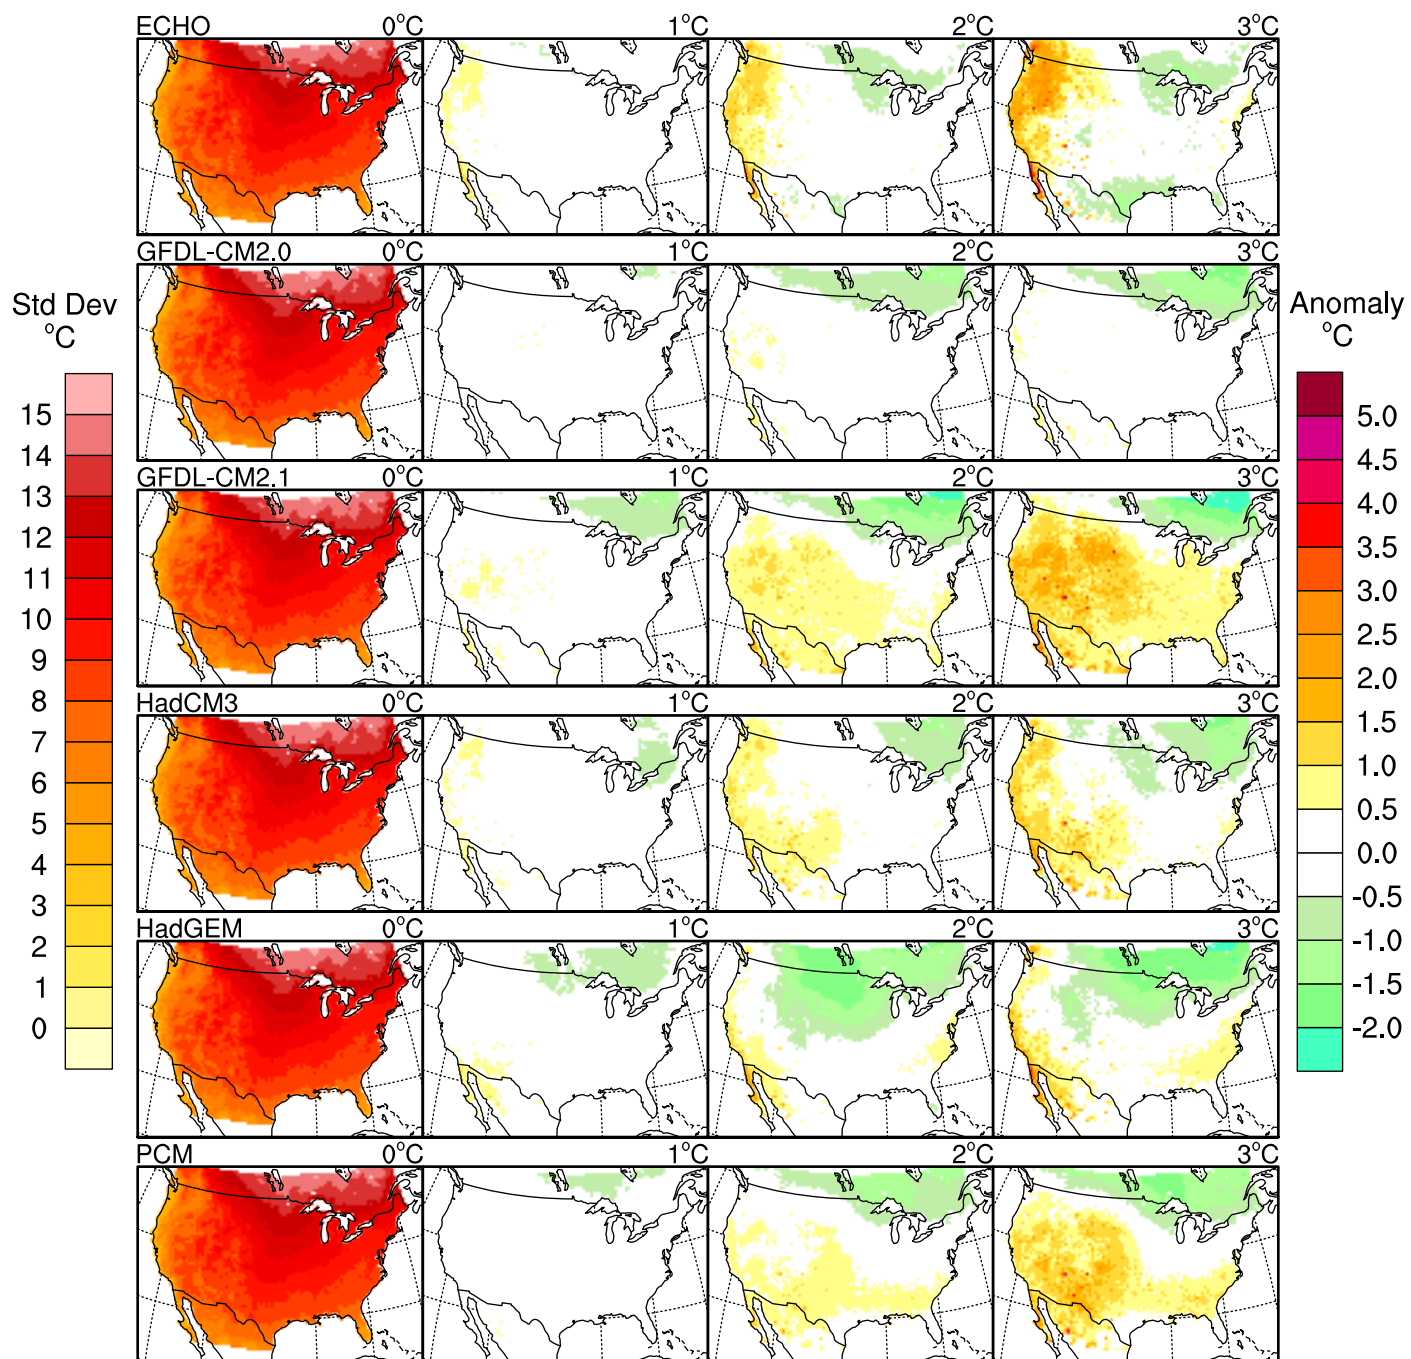
Fall (SON) Tmin Standard Deviation



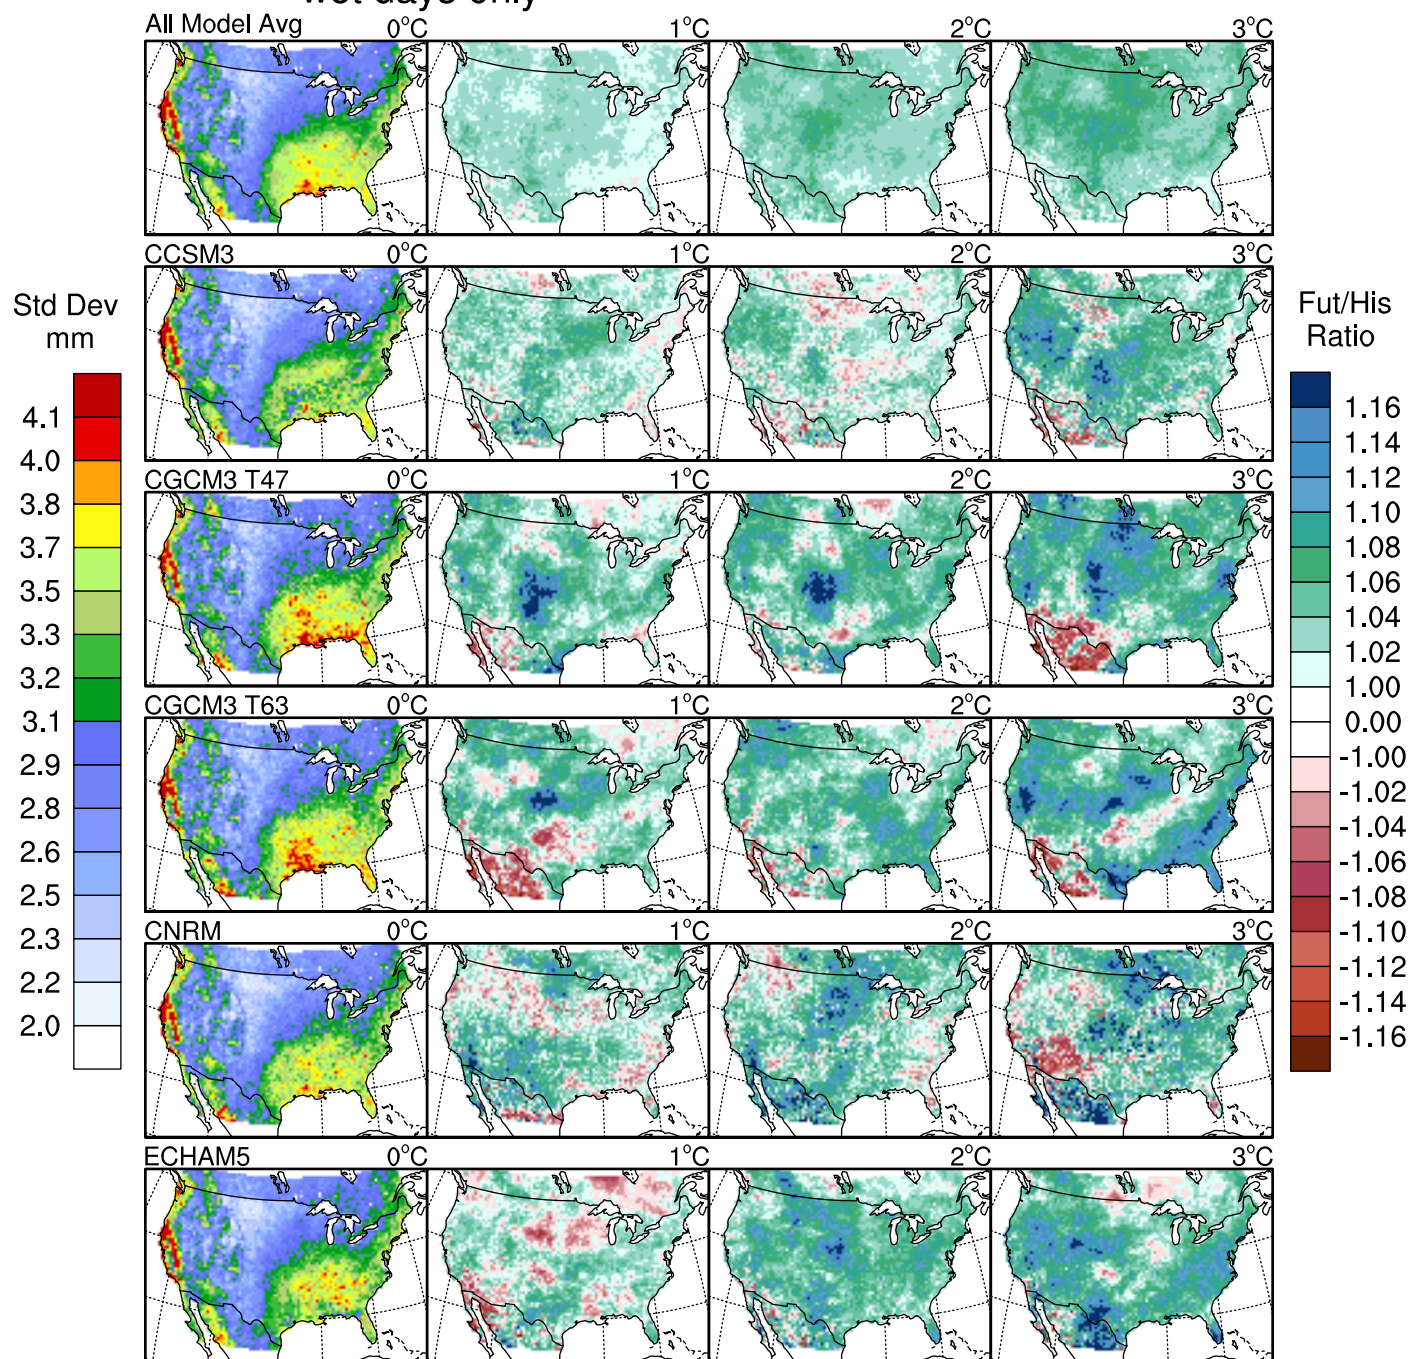
Annual Tmin Standard Deviation



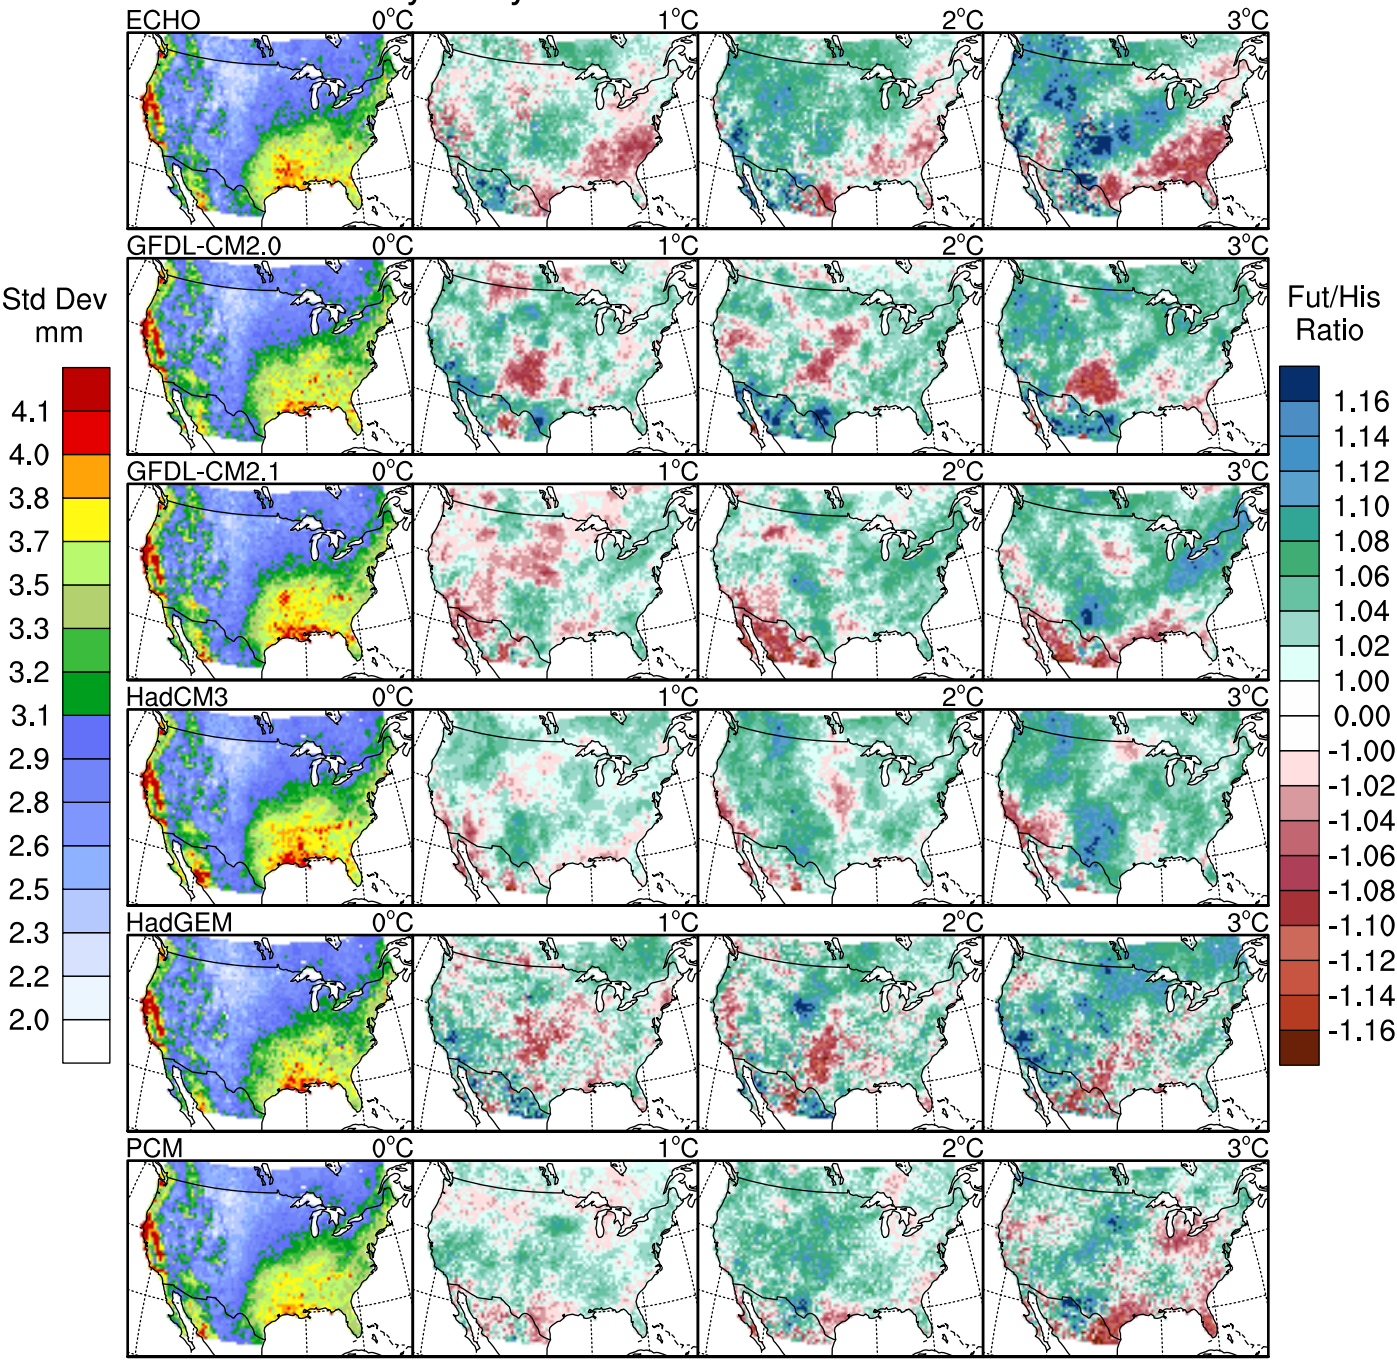
Annual Tmin Standard Deviation



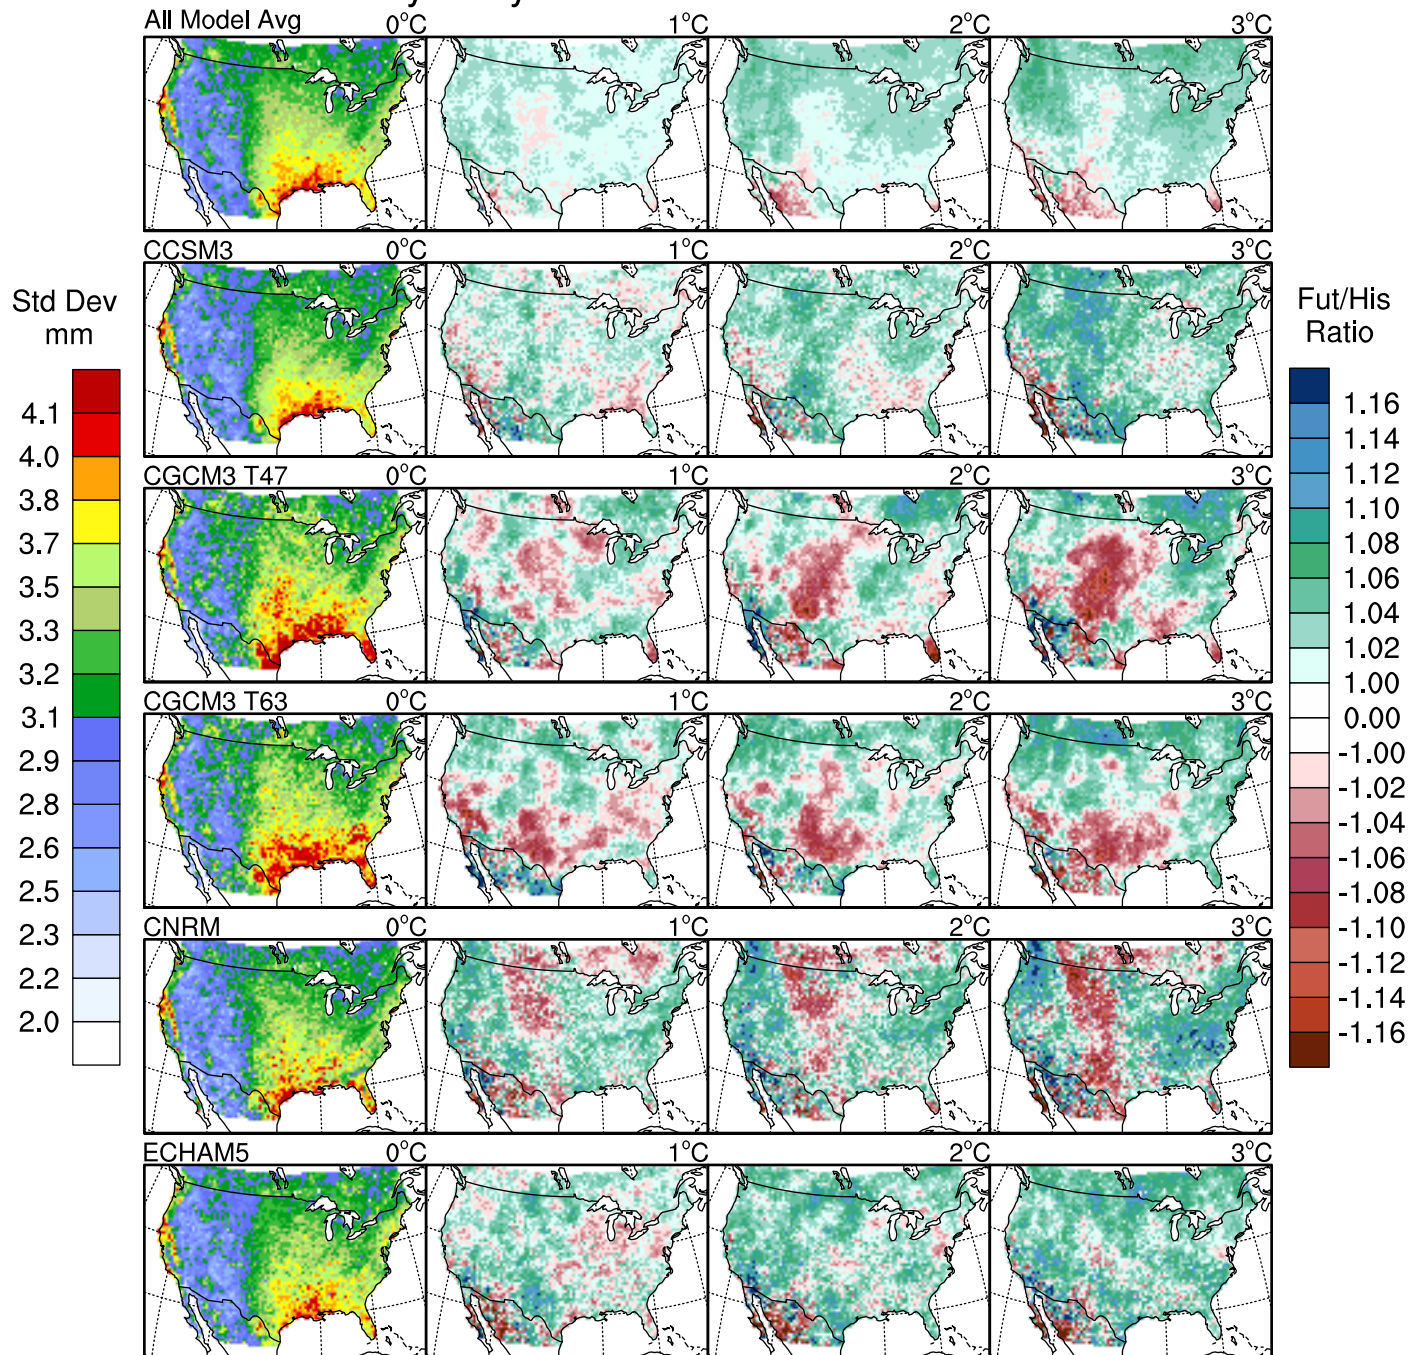
Winter (DJF) Cumulative Pr Standard Deviation wet days only



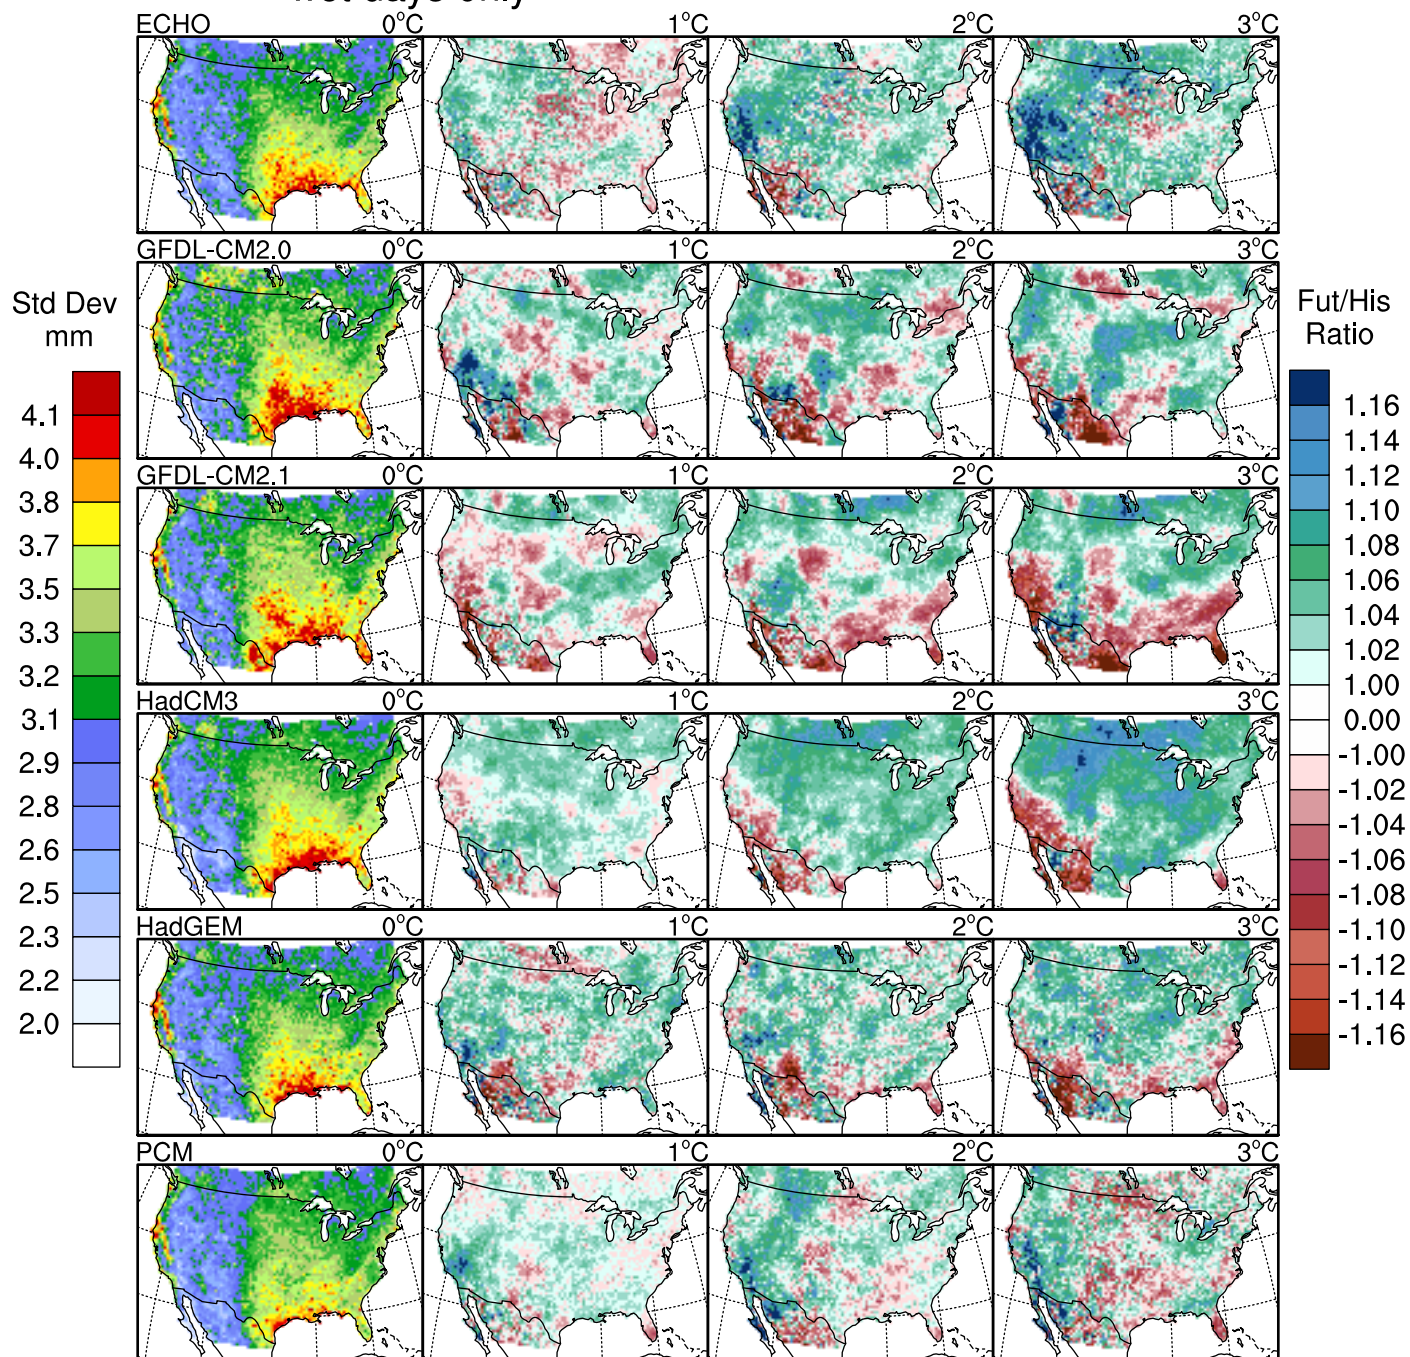
Winter (DJF) Cumulative Pr Standard Deviation
wet days only



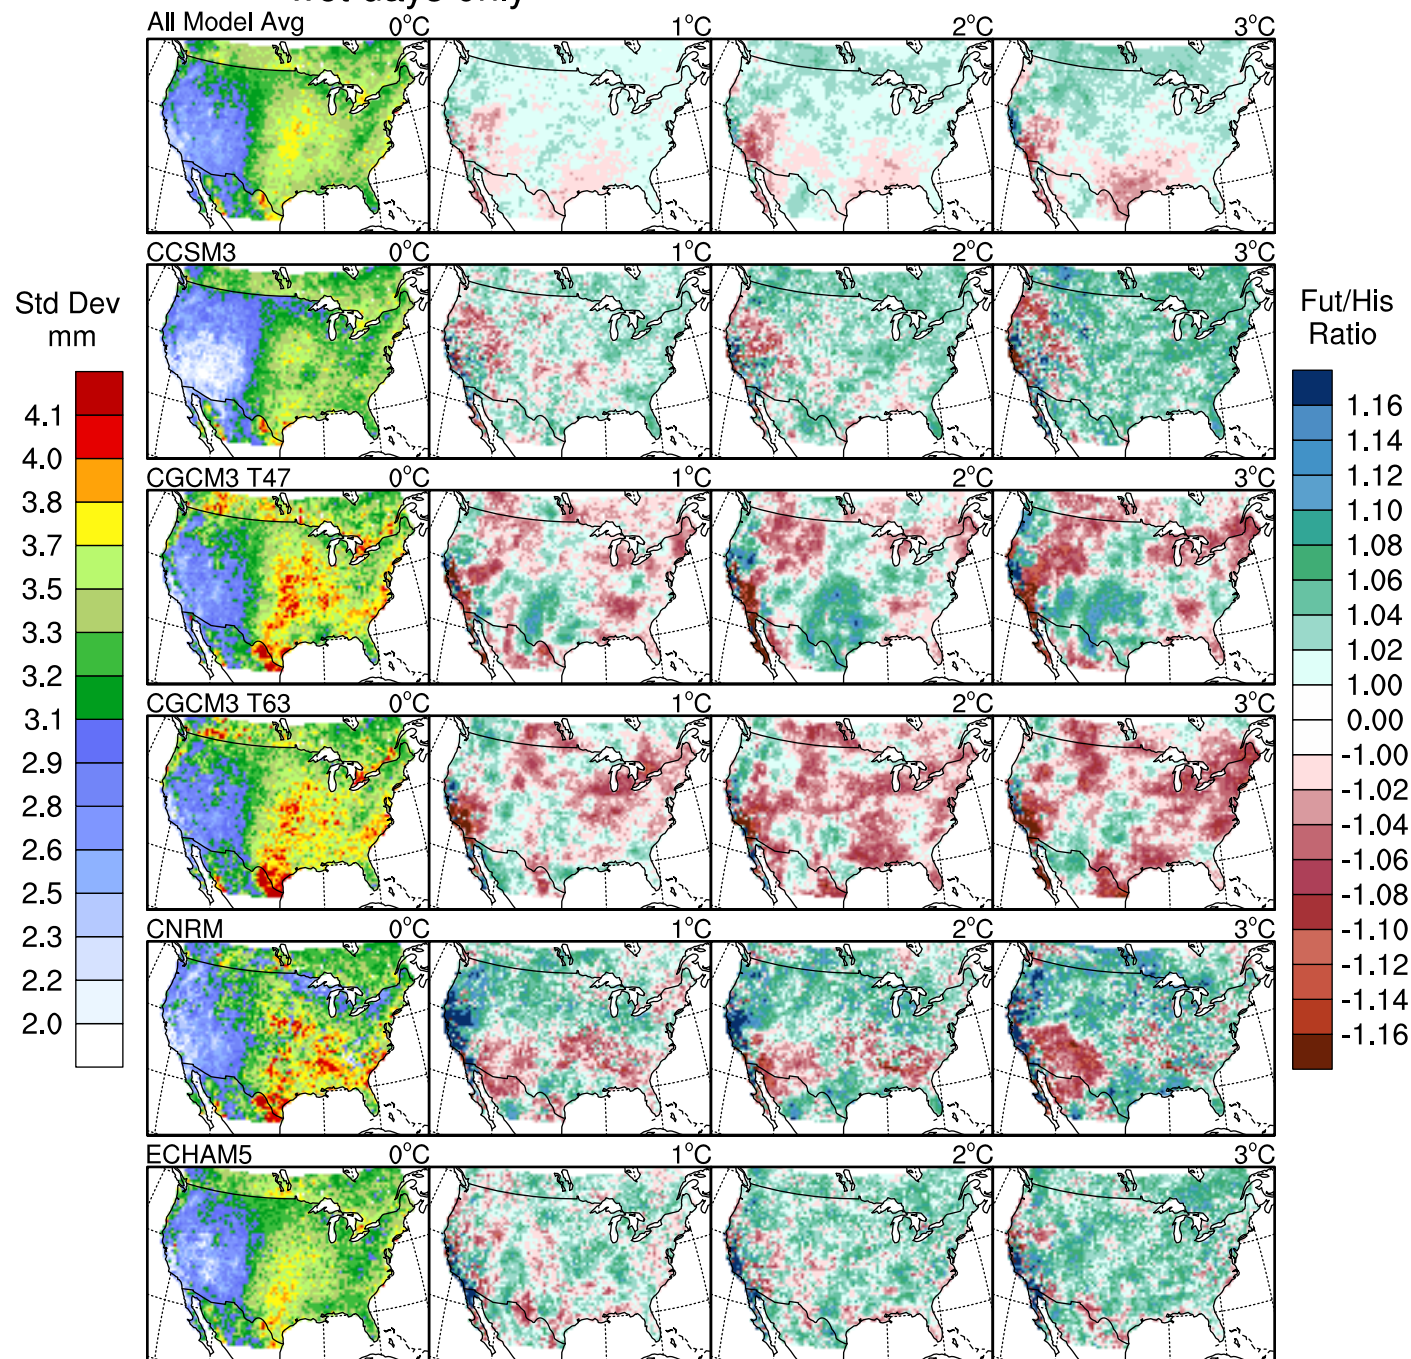
Spring (MAM) Cumulative Pr Standard Deviation wet days only



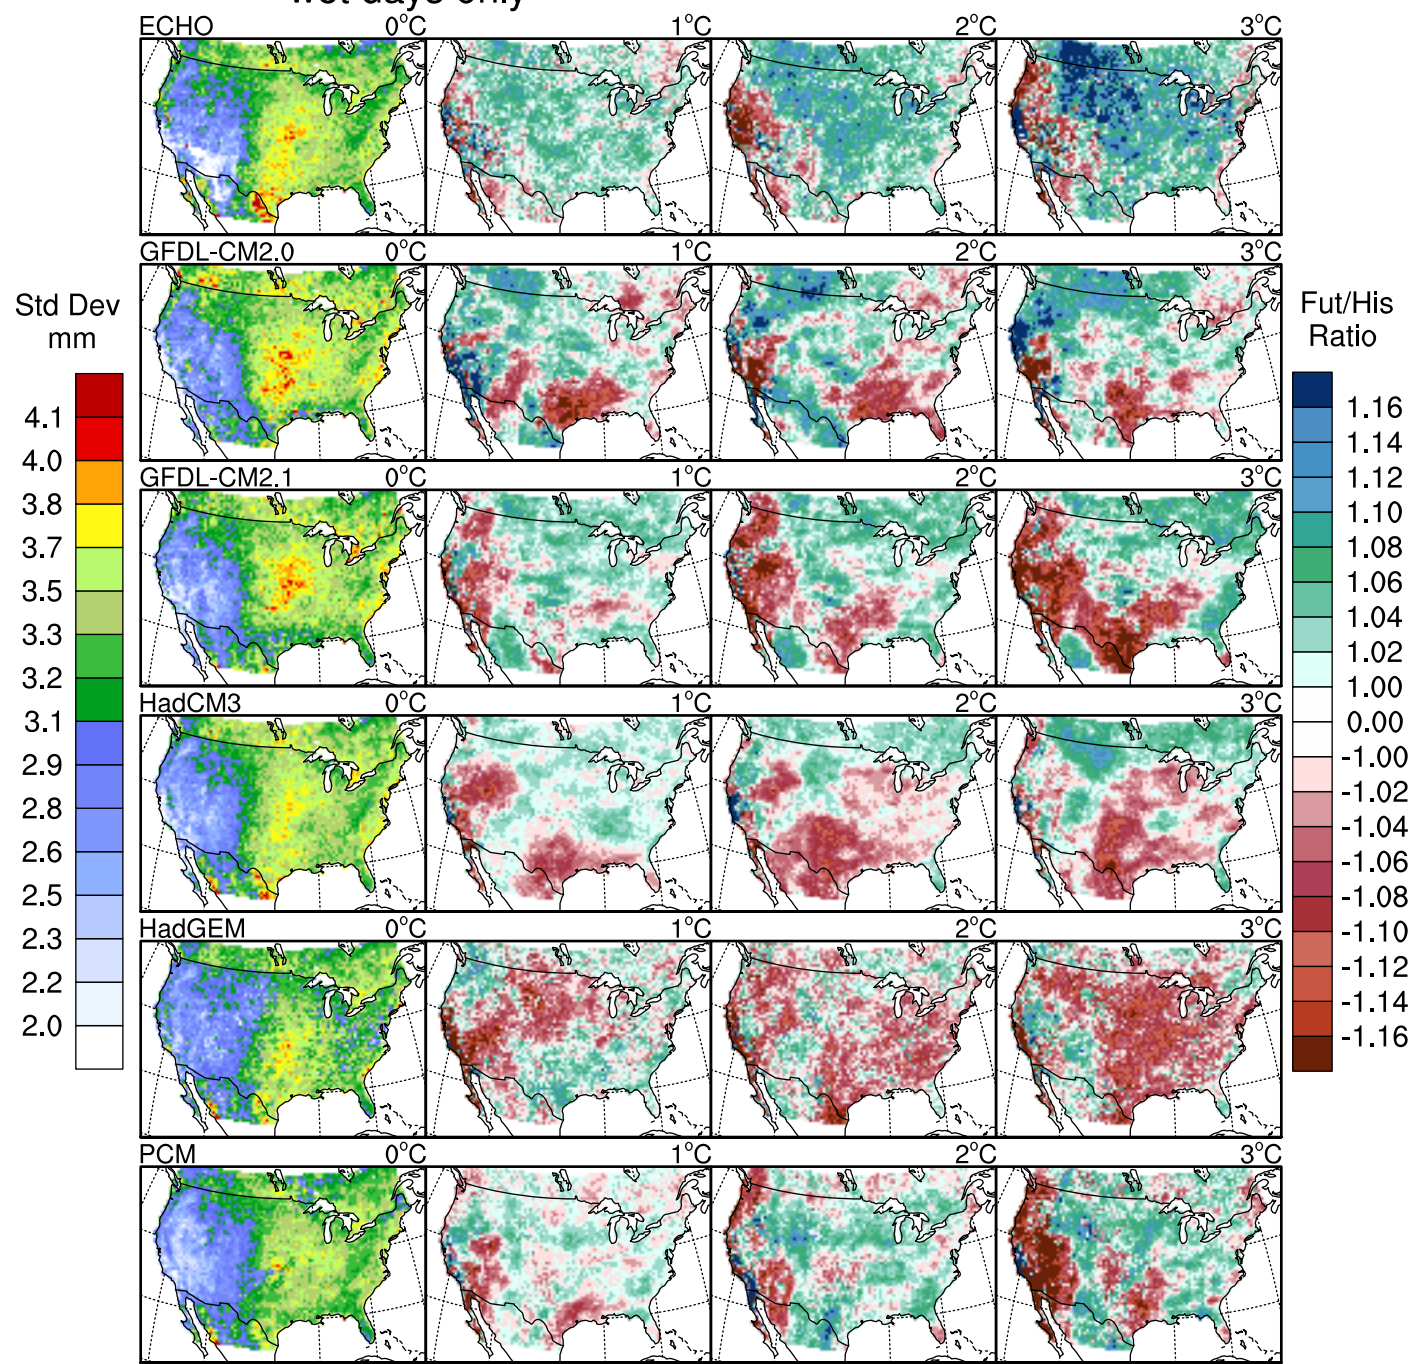
Spring (MAM) Cumulative Pr Standard Deviation wet days only



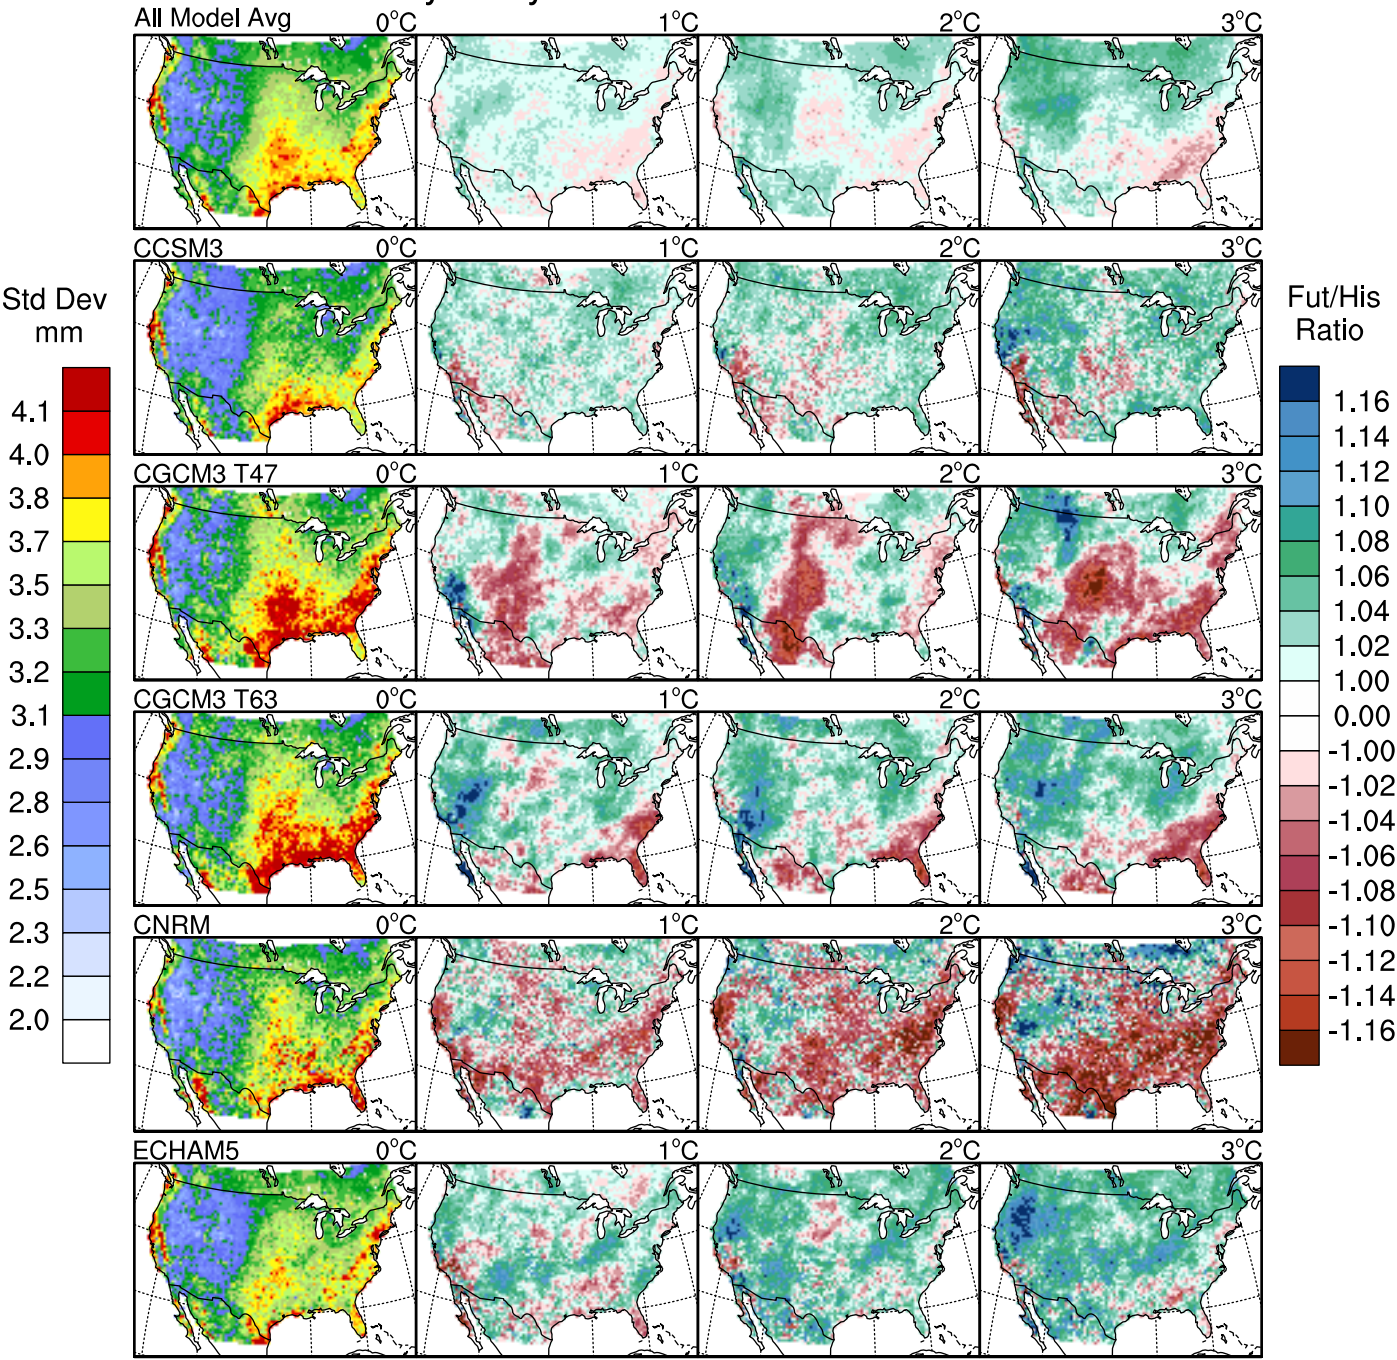
Summer (JJA) Cumulative Pr Standard Deviation
wet days only



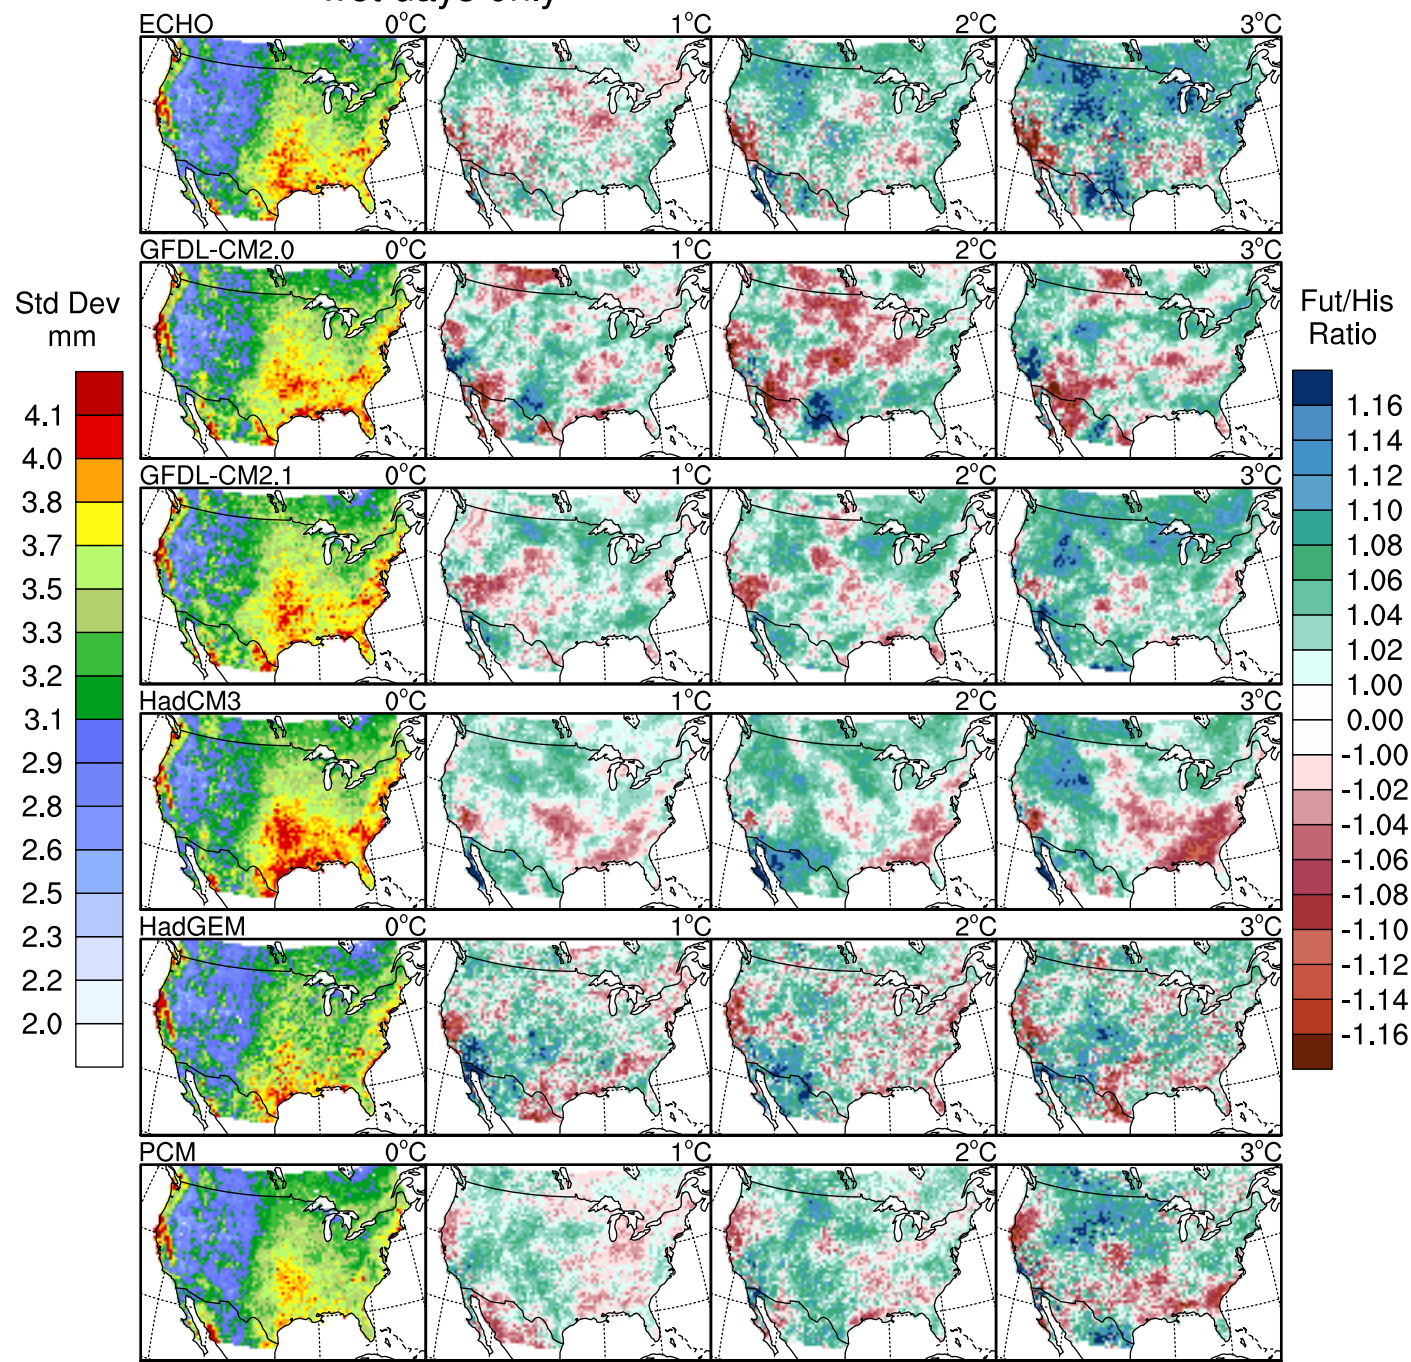
Summer (JJA) Cumulative Pr Standard Deviation
wet days only



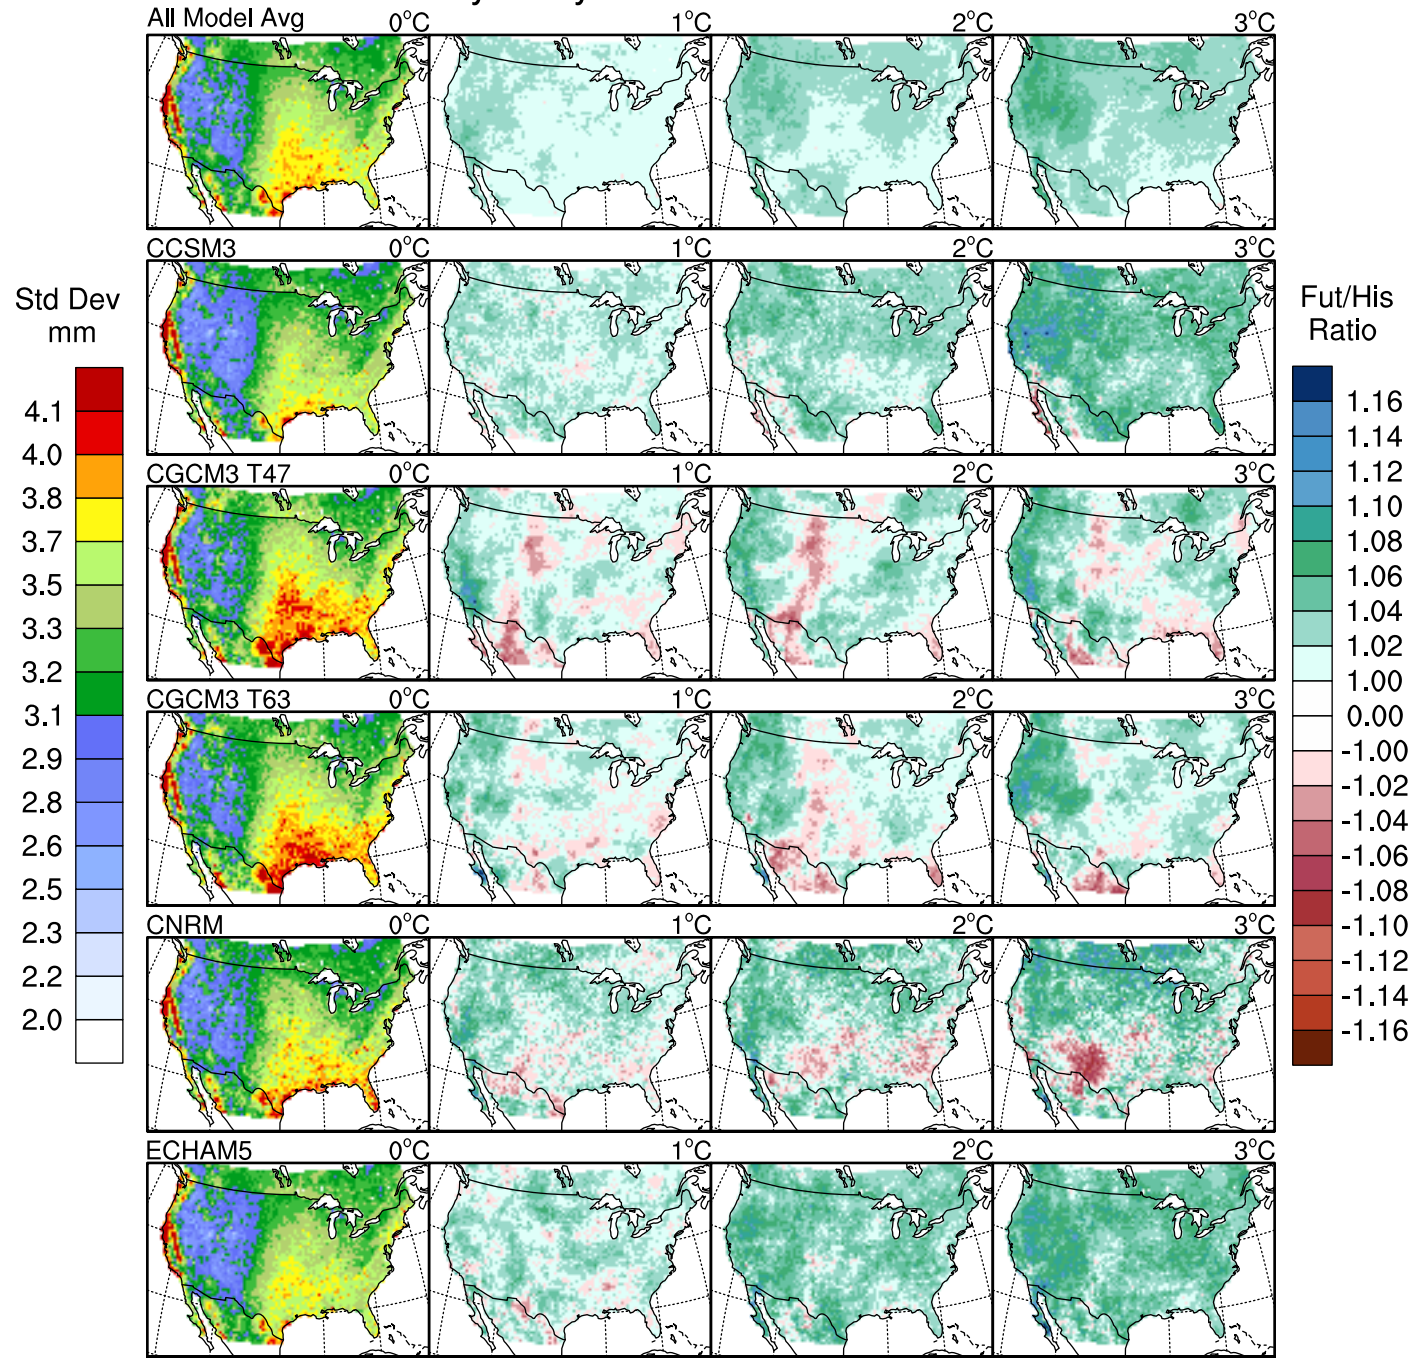
Fall (SON) Cumulative Pr Standard Deviation
wet days only




Fall (SON) Cumulative Pr Standard Deviation
wet days only



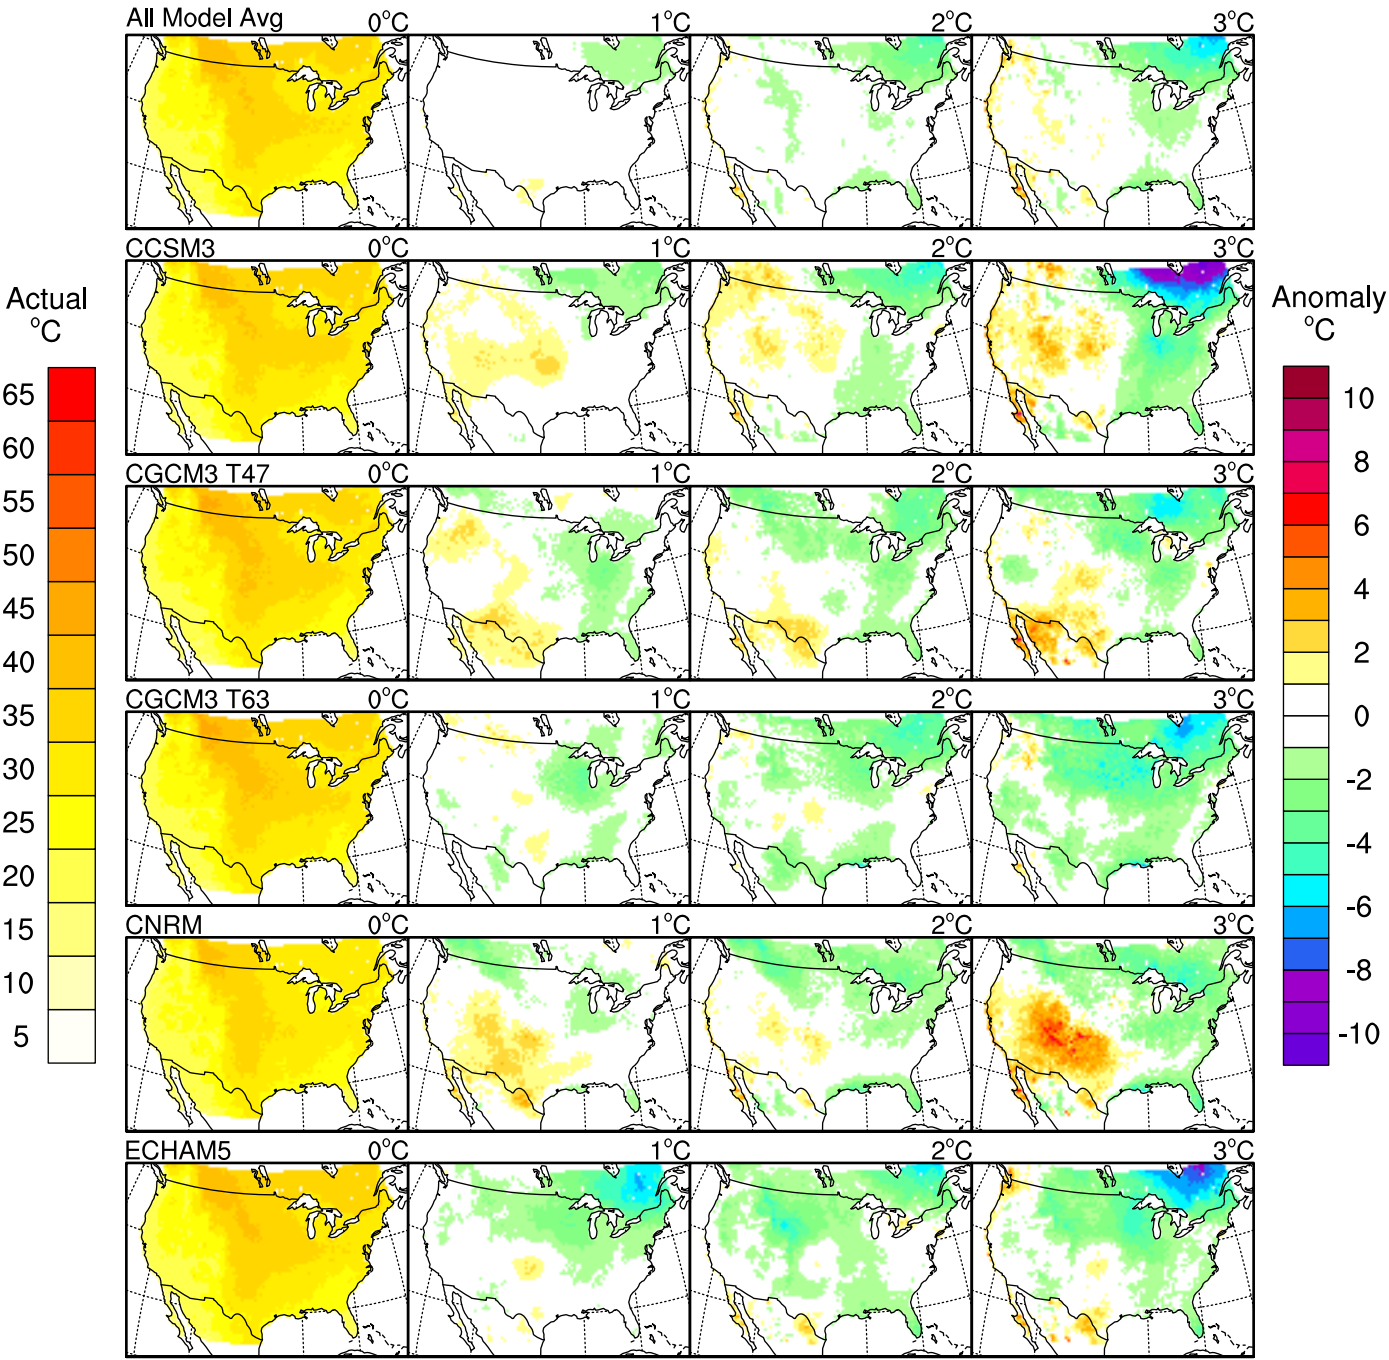
Annual Cumulative Pr Standard Deviation
wet days only




Annual Cumulative Pr Standard Deviation
wet days only



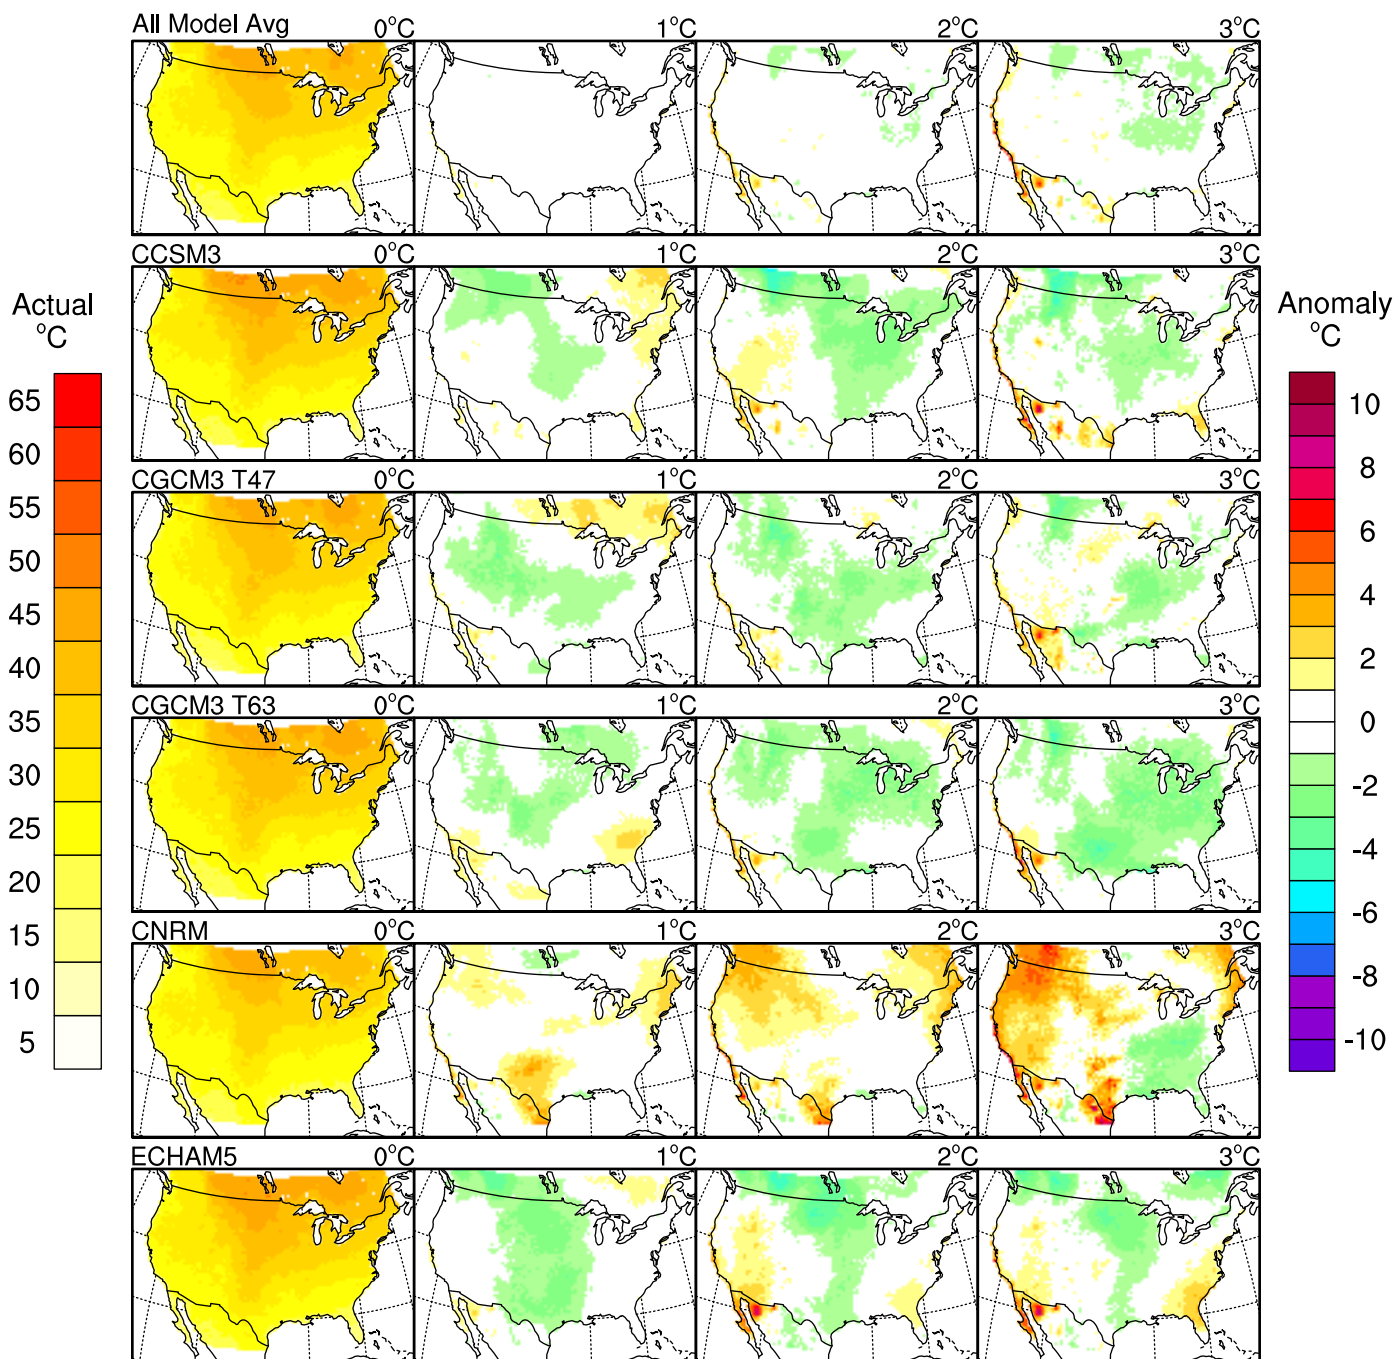
Range in Daily Winter (DJF) Tmax



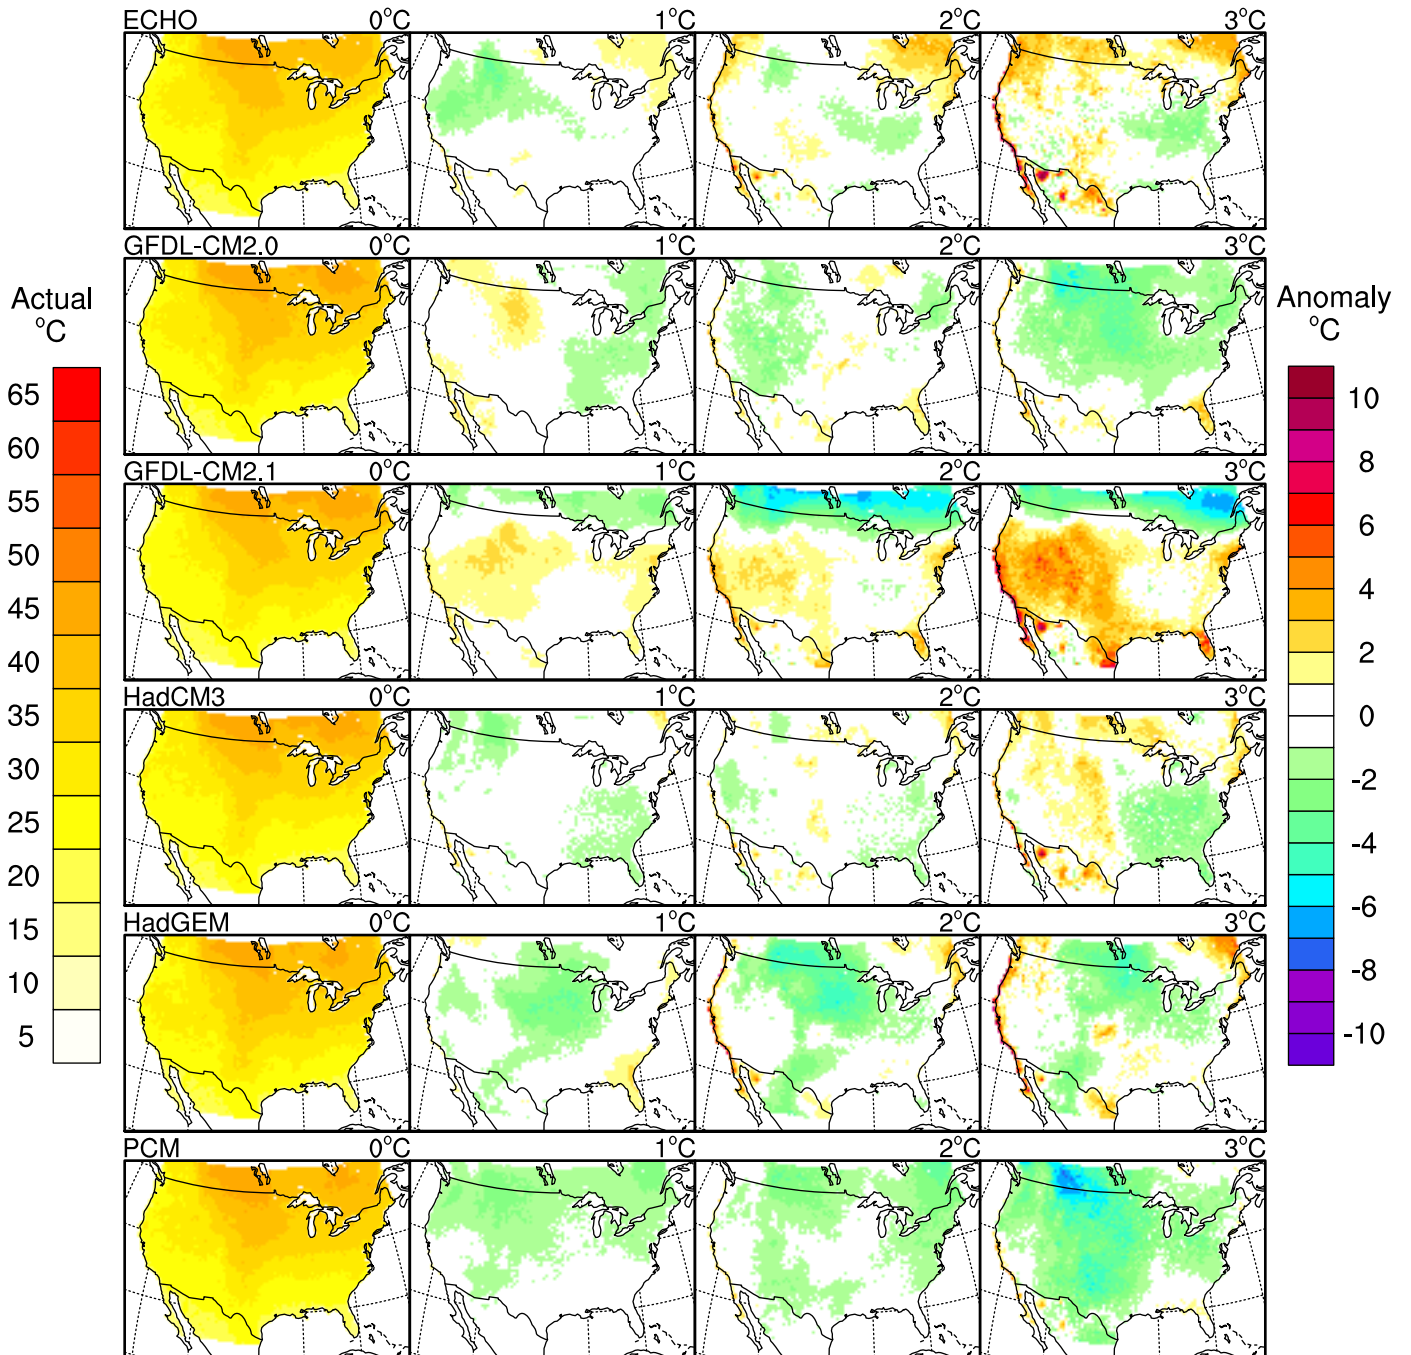
Range in Daily Winter (DJF) Tmax




Range in Daily Spring (MAM) Tmax



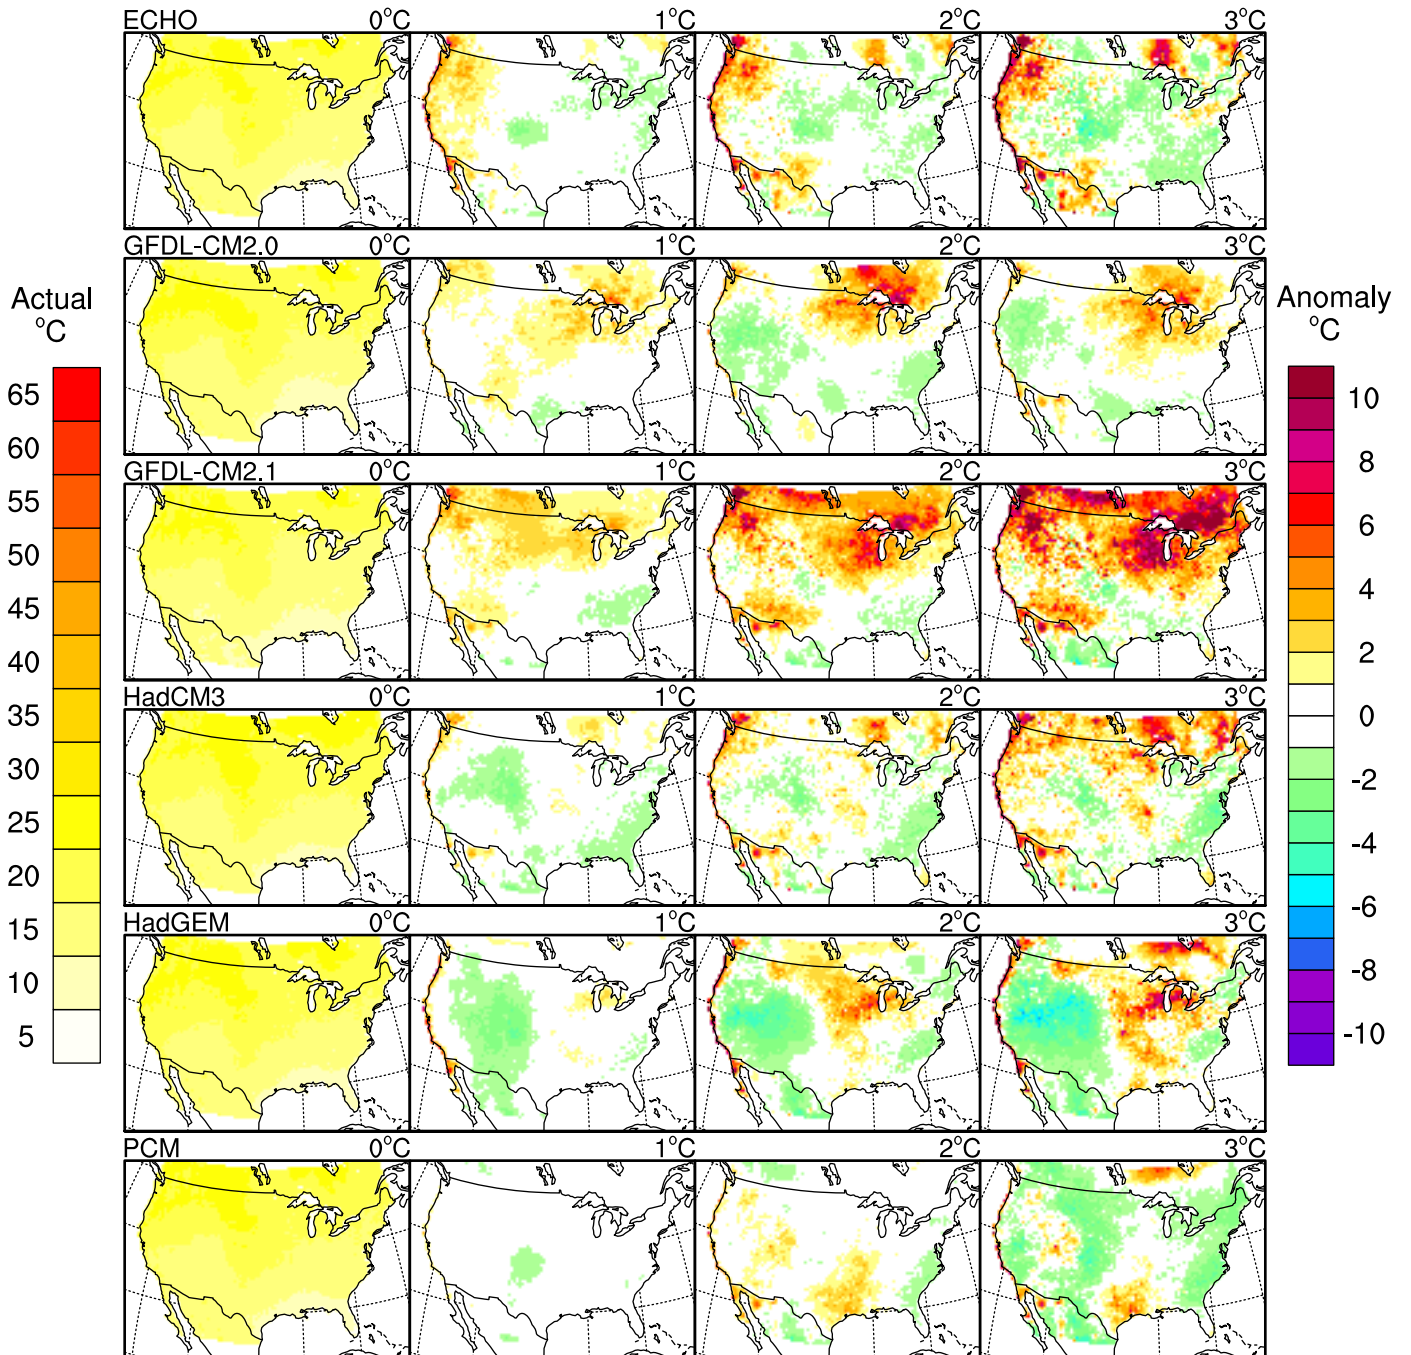
Range in Daily Spring (MAM) Tmax



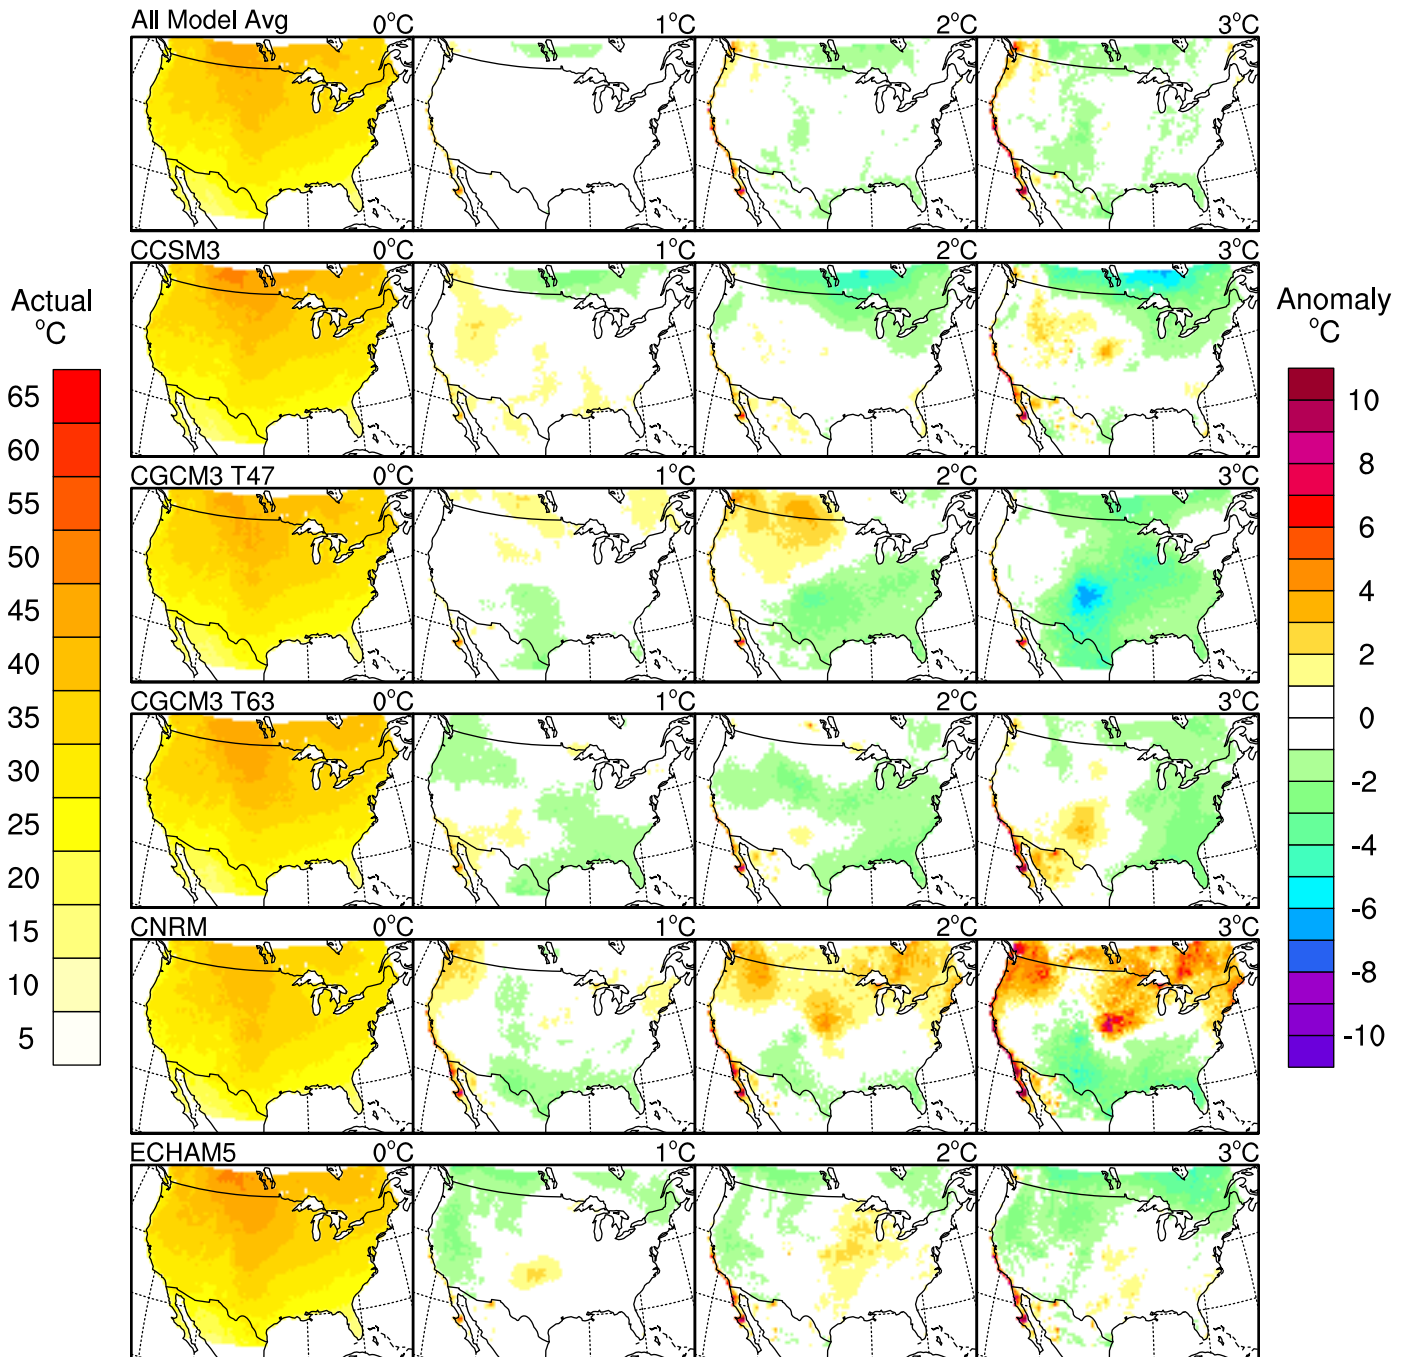
Range in Daily Summer (JJA) Tmax



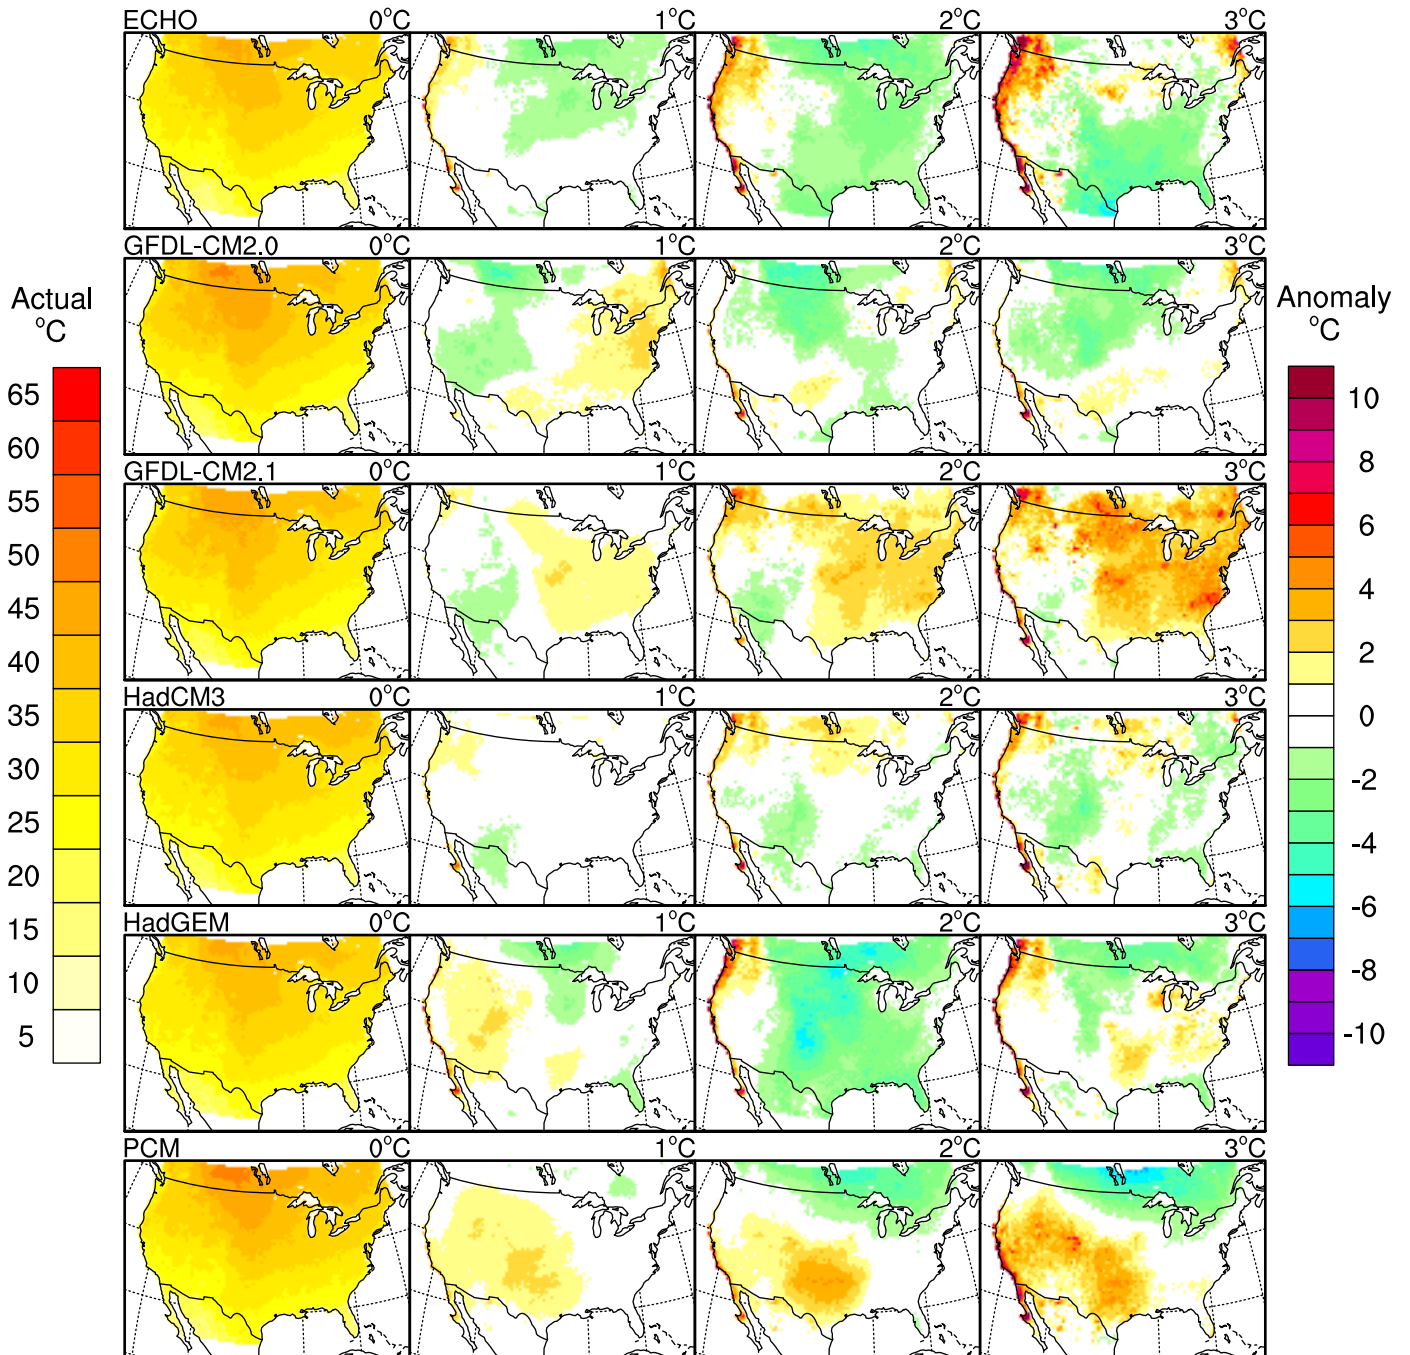
Range in Daily Summer (JJA) Tmax




Range in Daily Fall (SON) Tmax



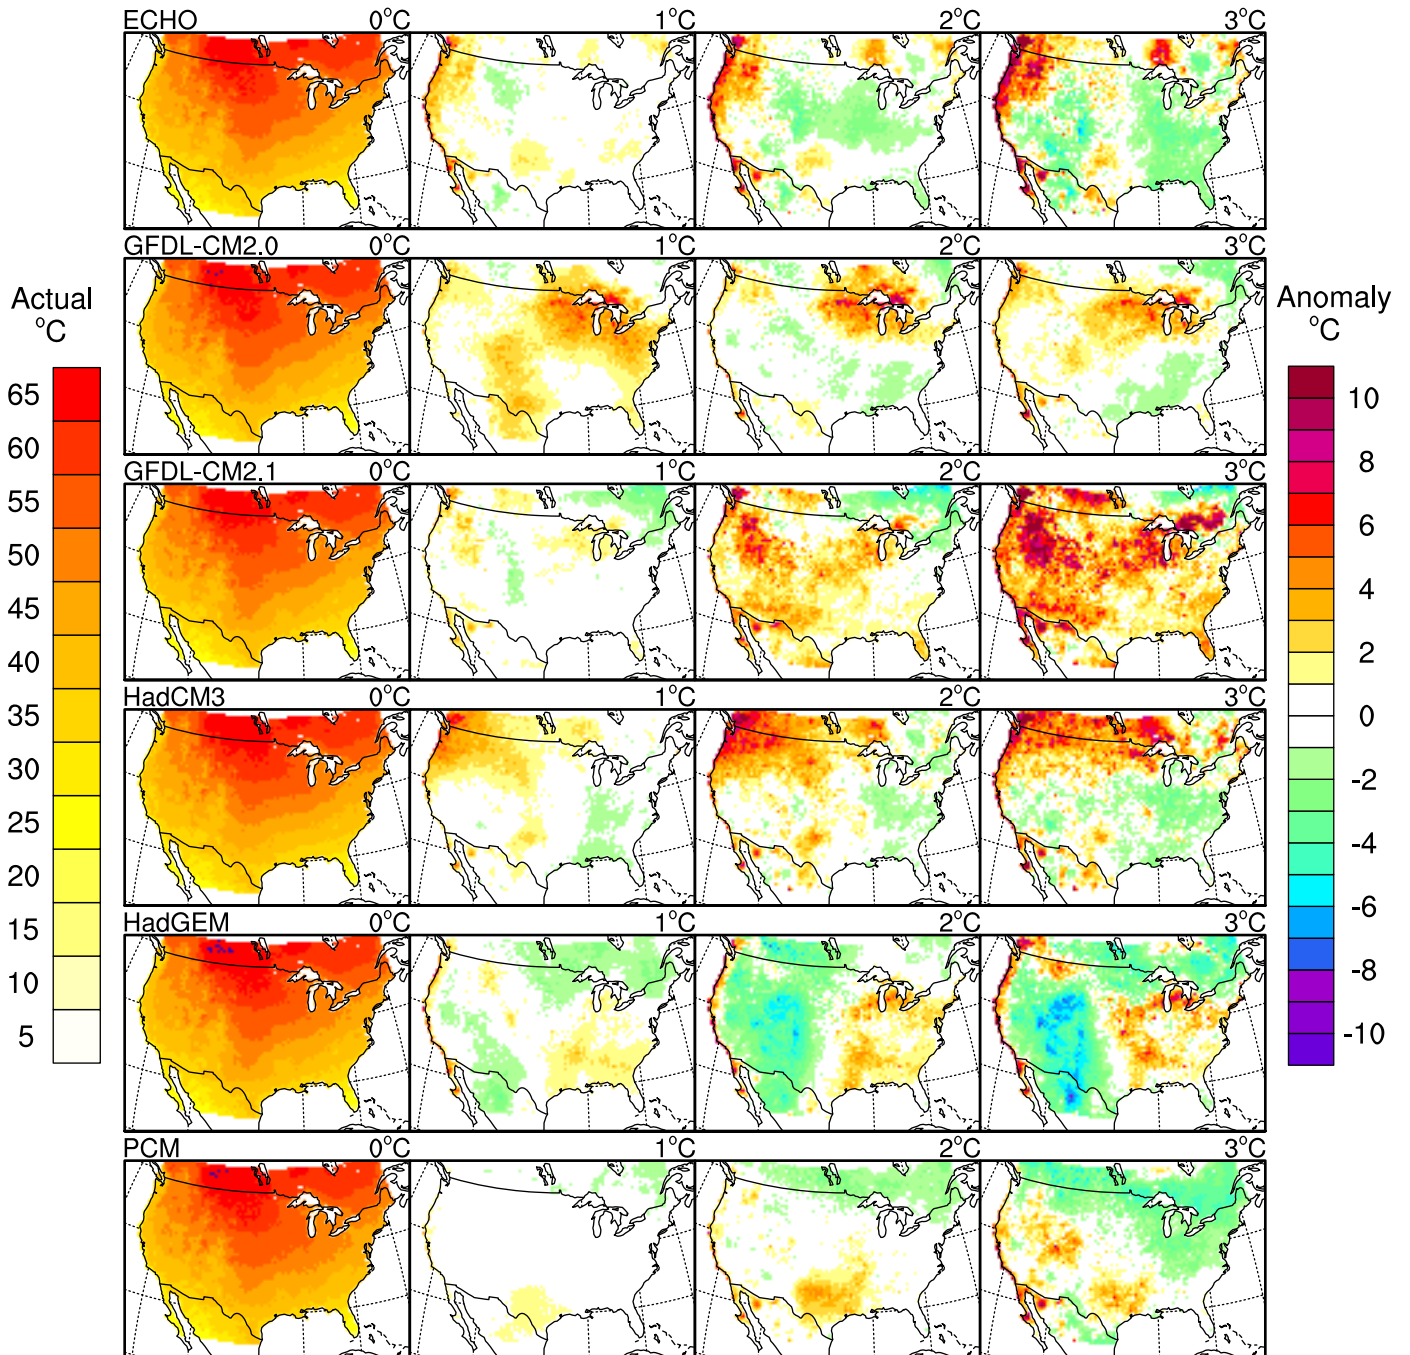
Range in Daily Fall (SON) Tmax



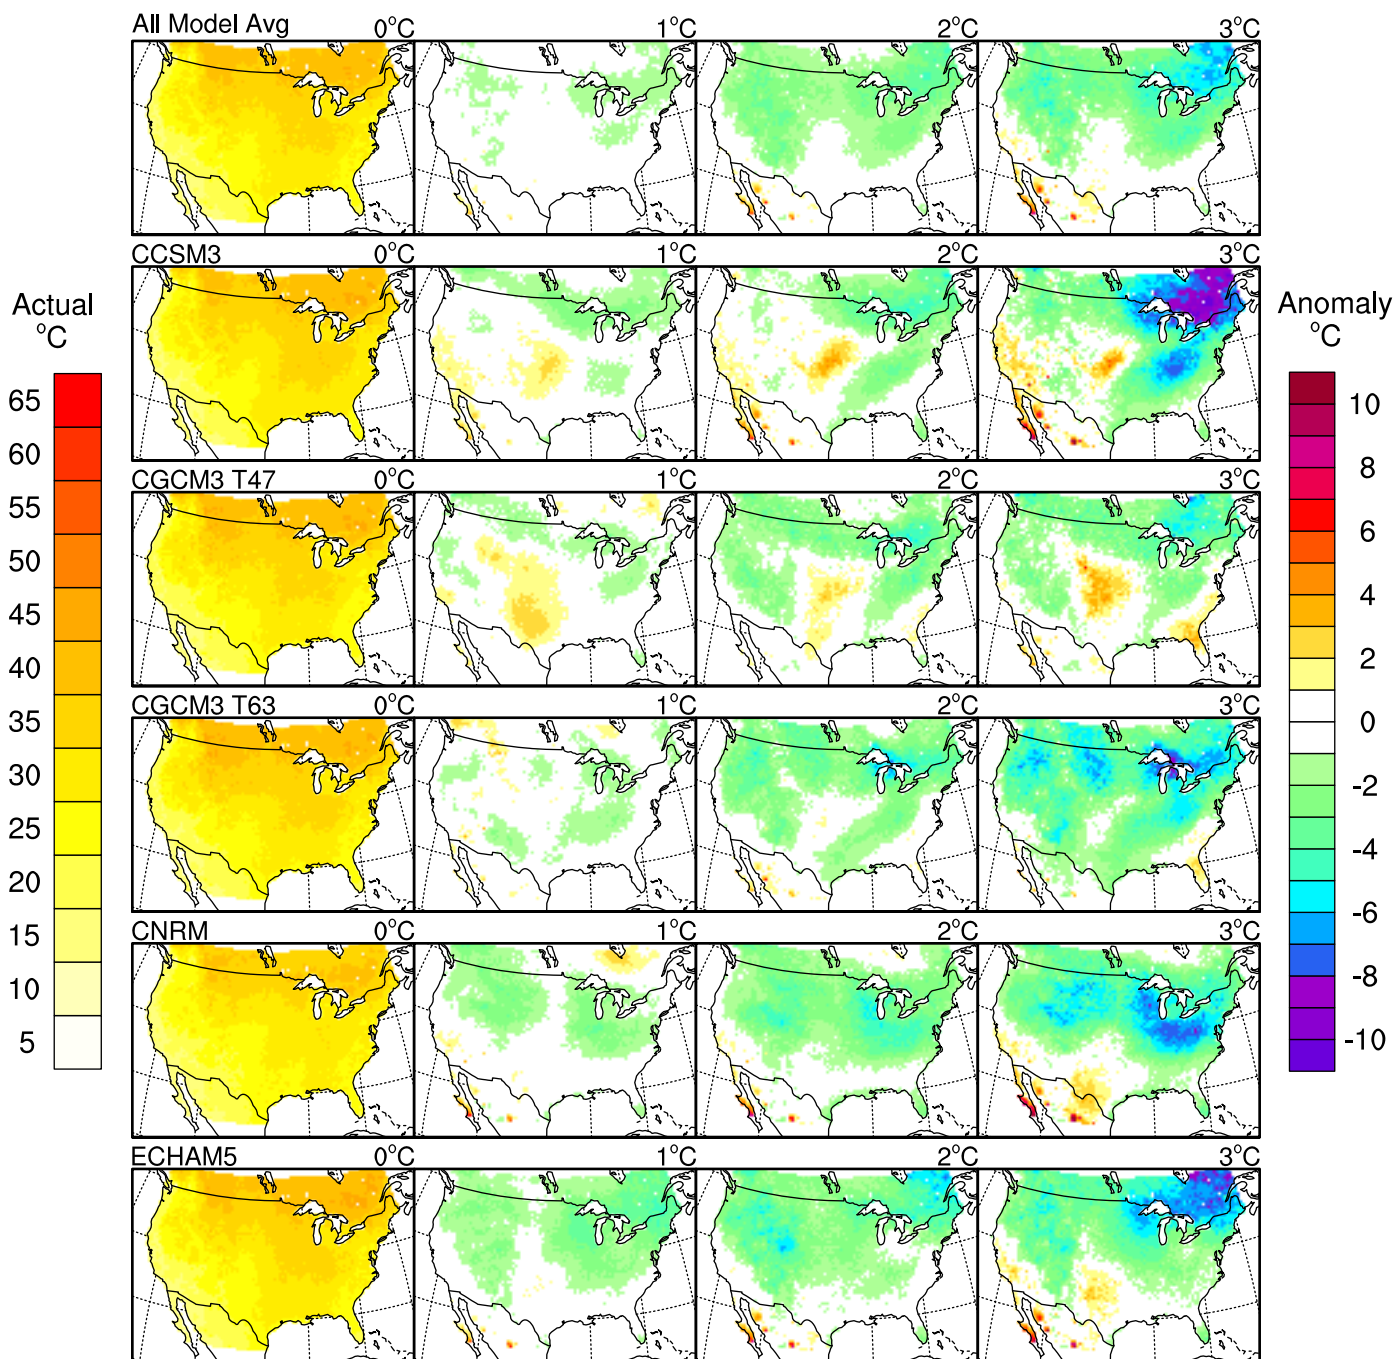
Range in Daily Annual Tmax



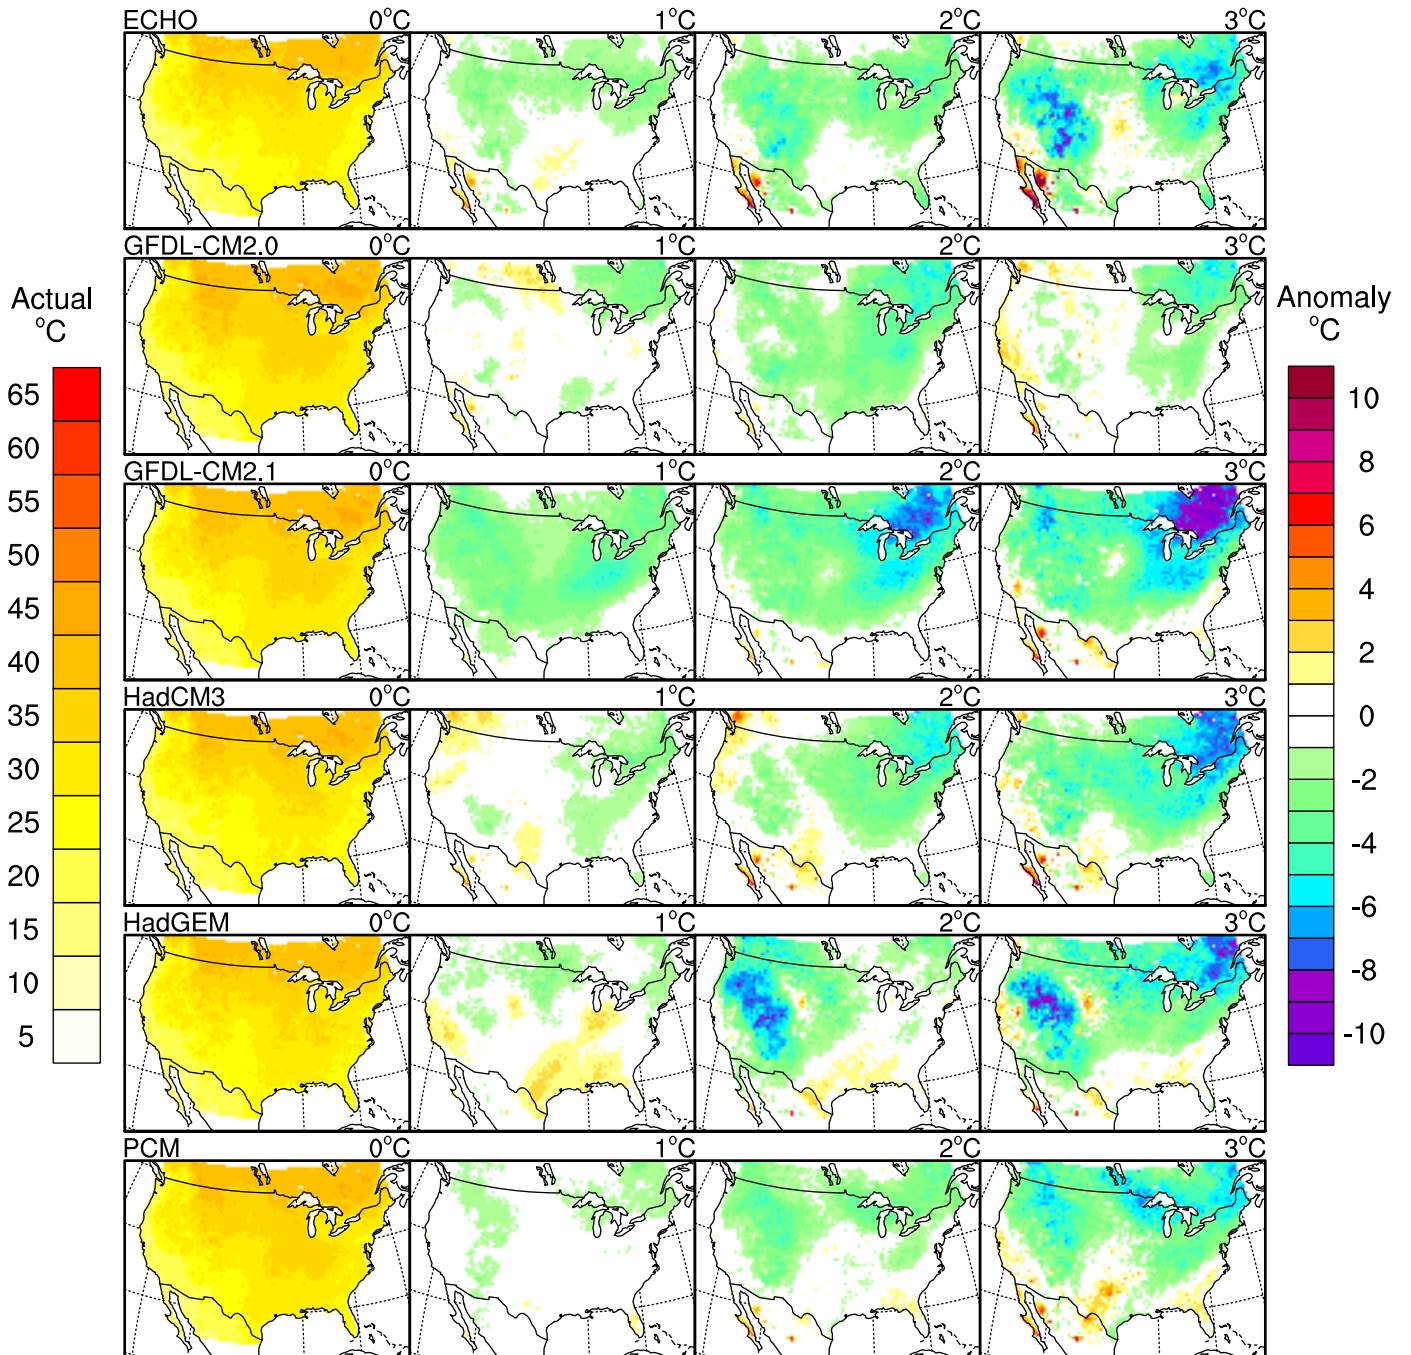
Range in Daily Annual Tmax




Range in Daily Winter (DJF) Tmin



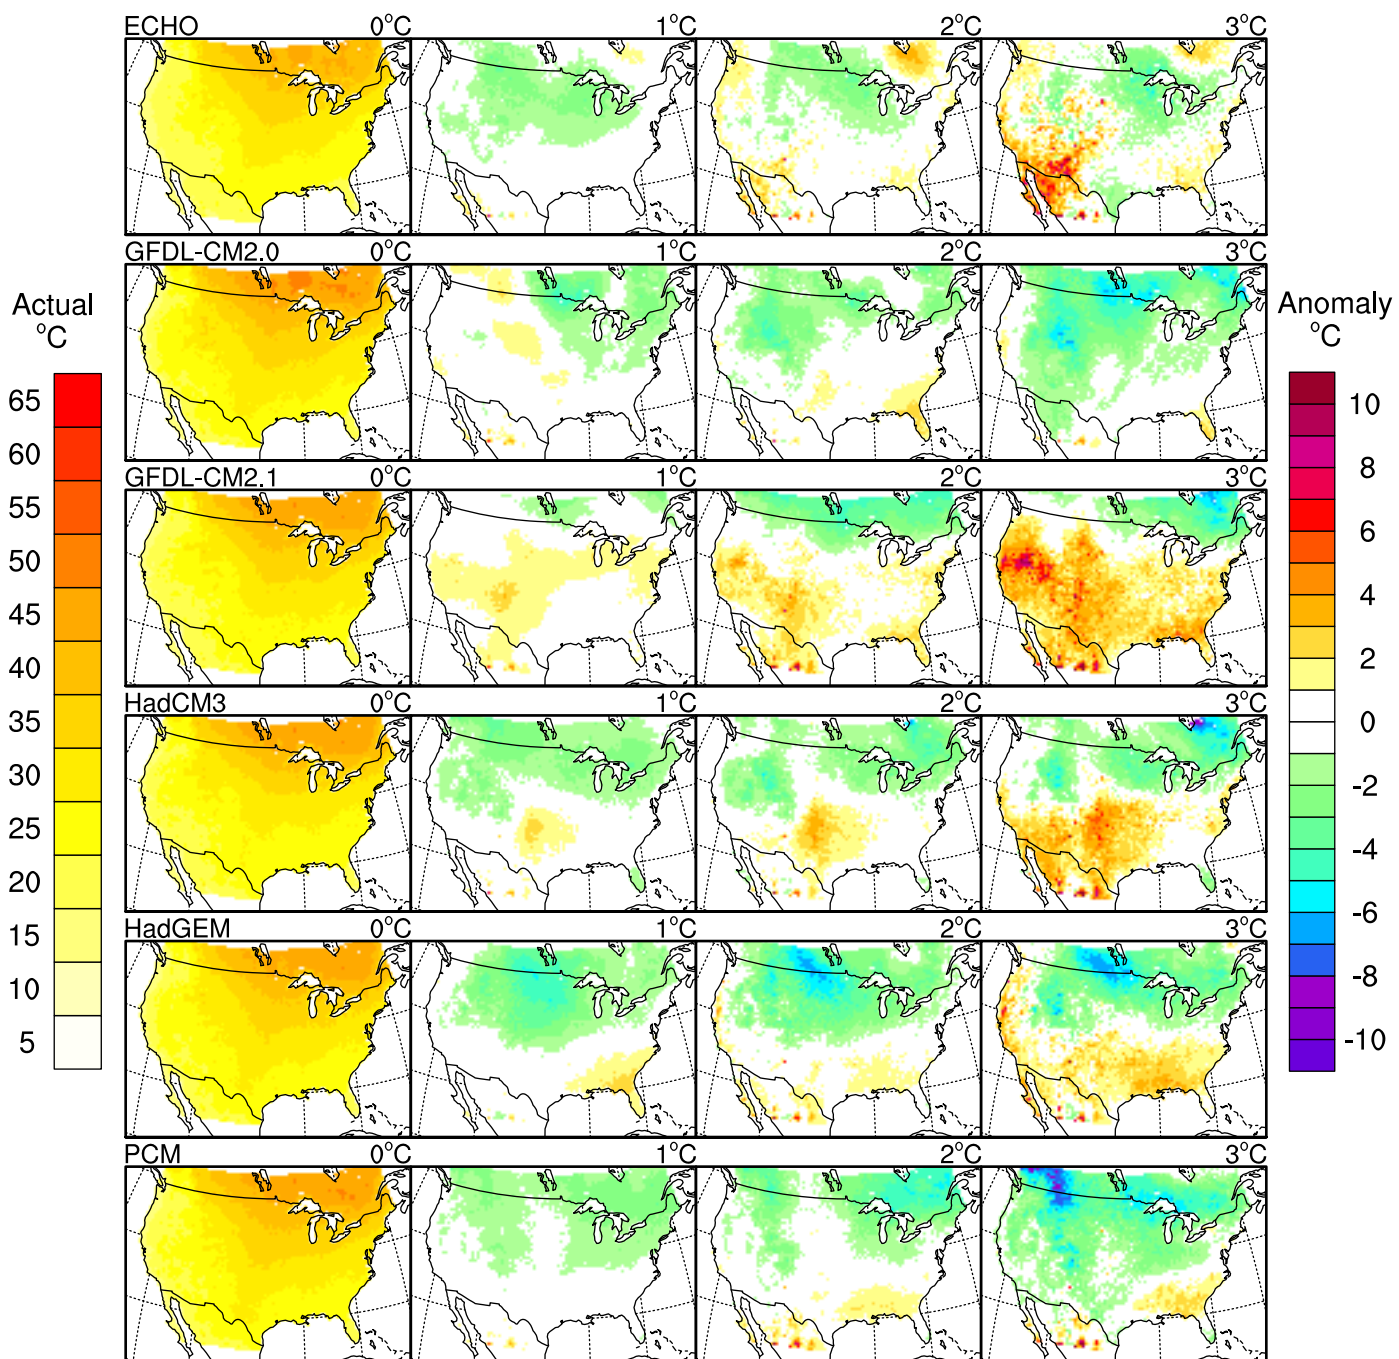
Range in Daily Winter (DJF) Tmin



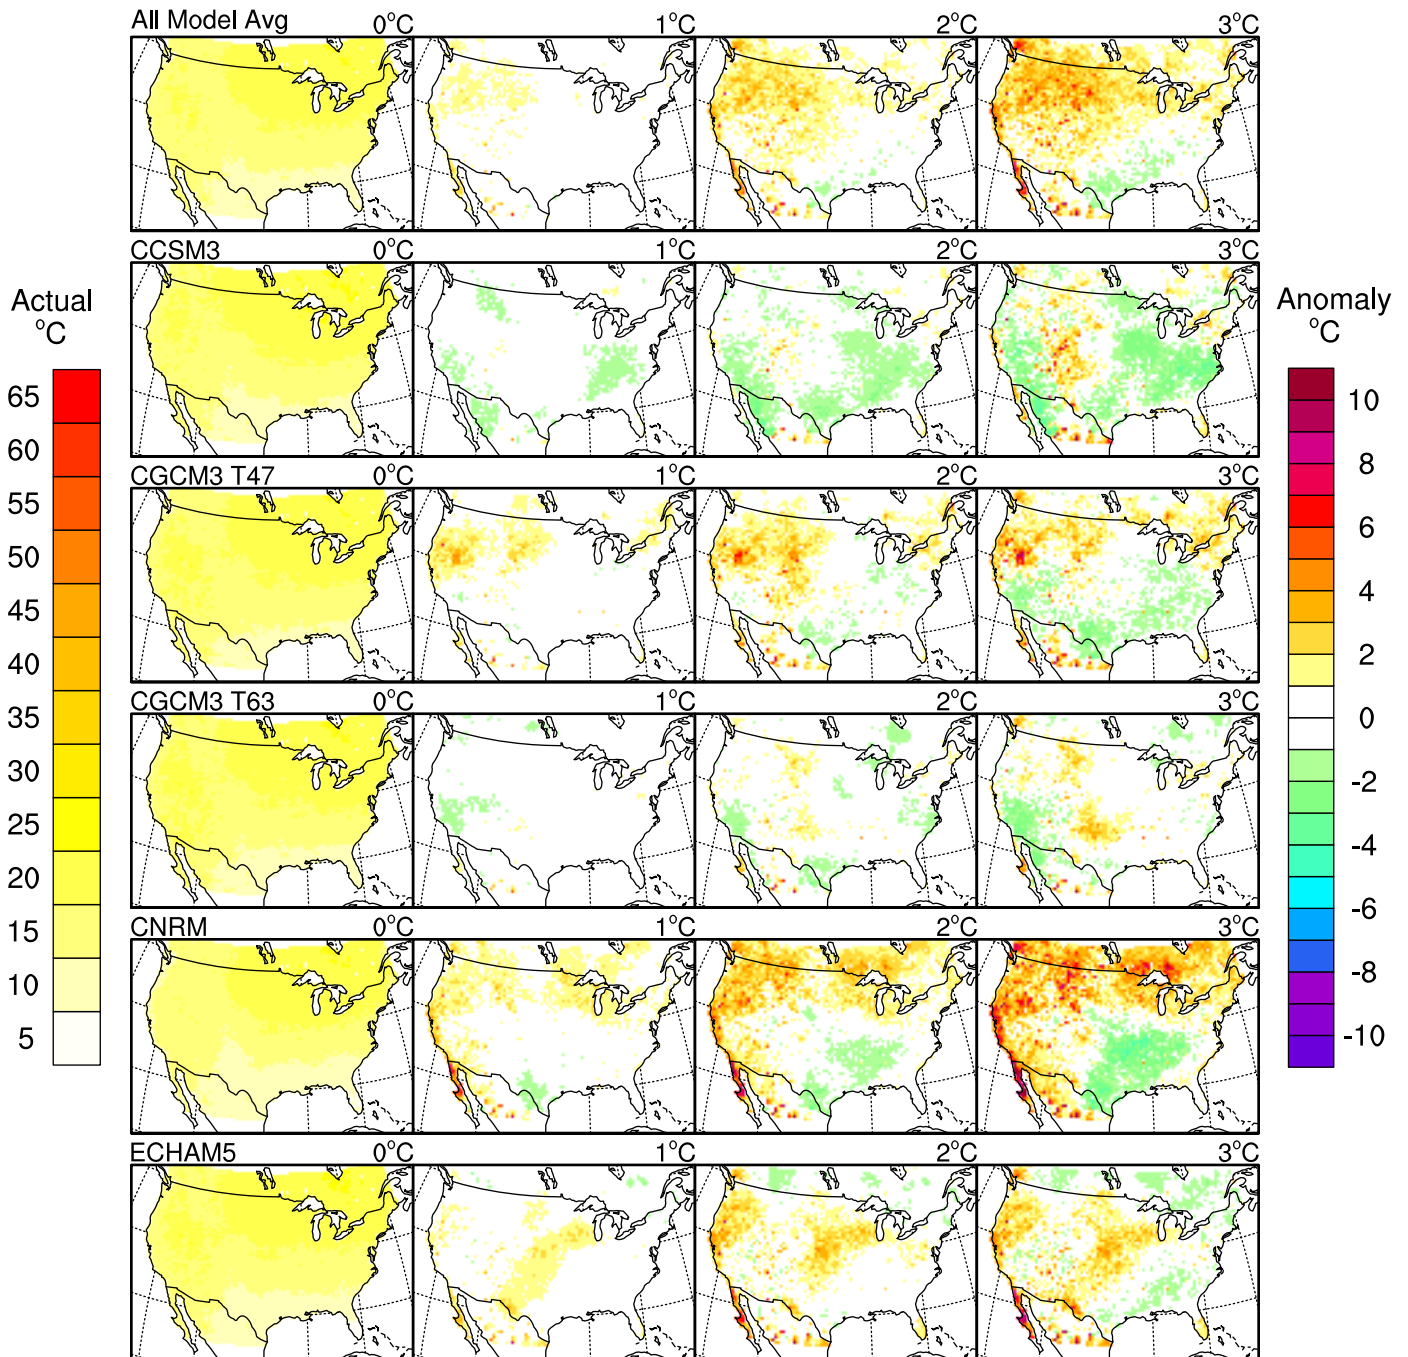
Range in Daily Spring (MAM) Tmin



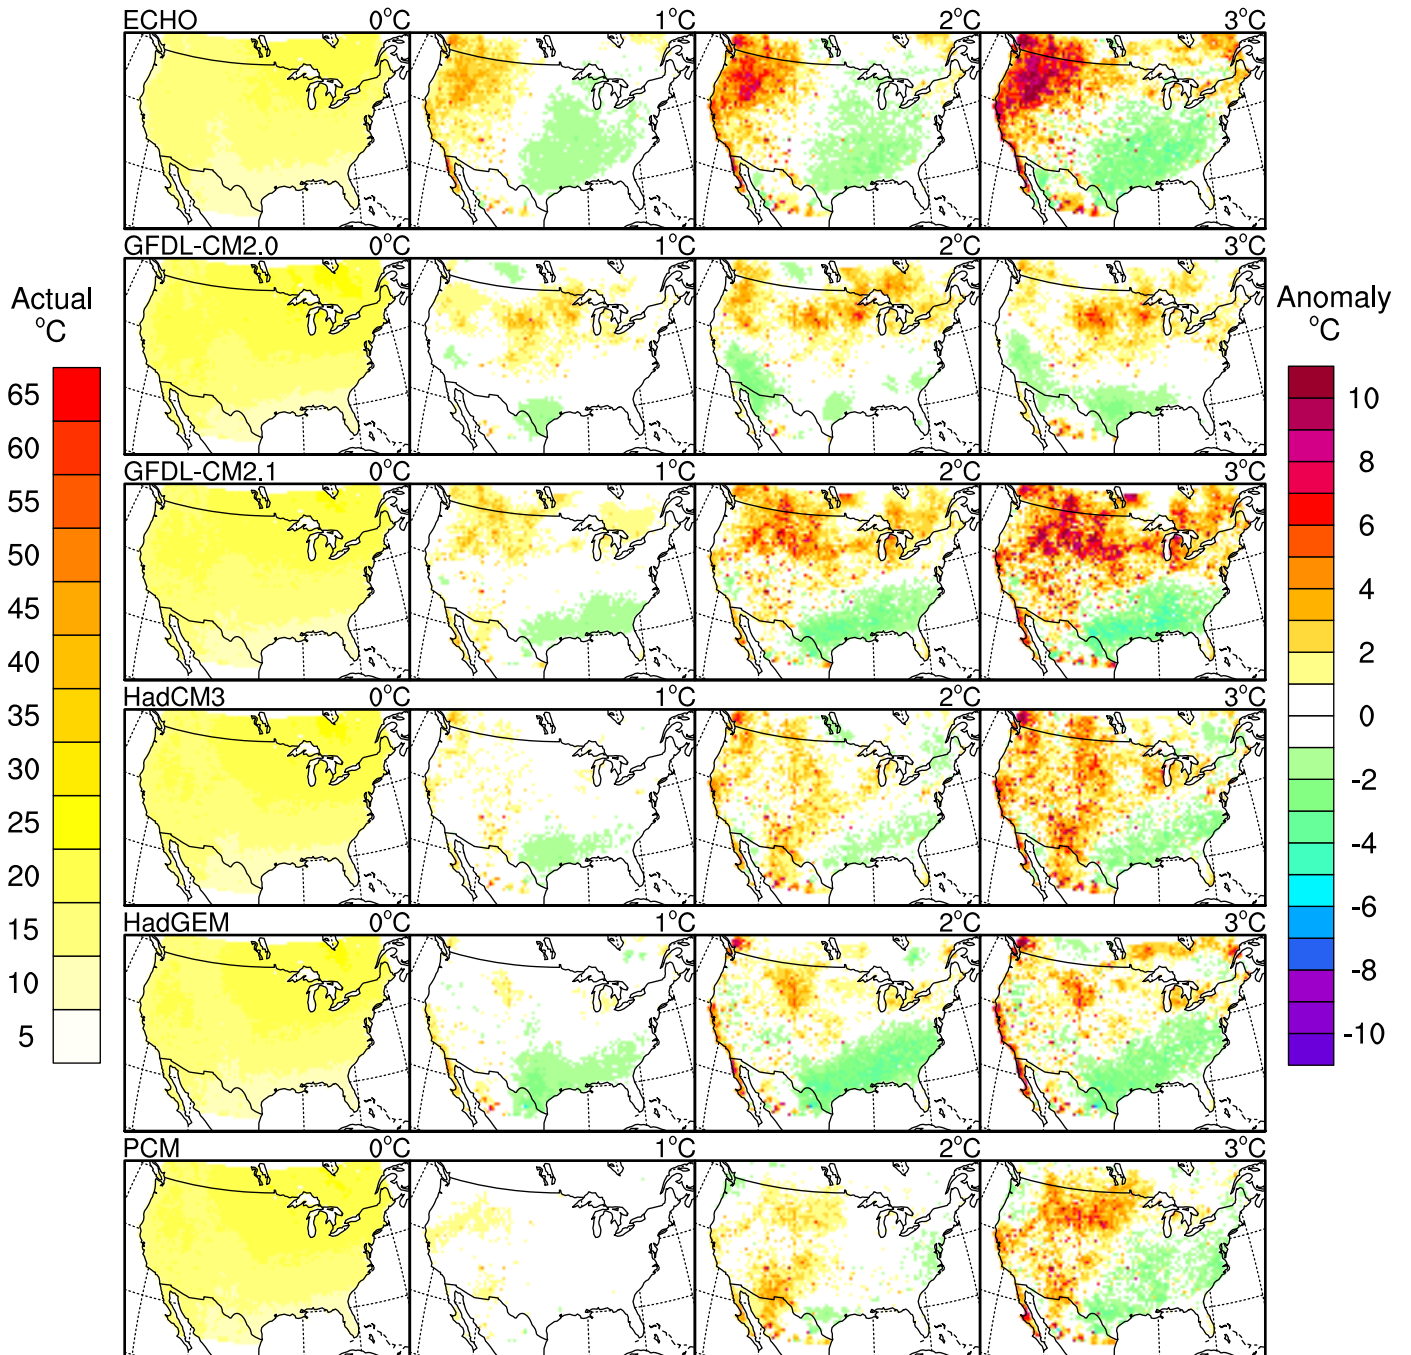
Range in Daily Spring (MAM) Tmin




Range in Daily Summer (JJA) Tmin




Range in Daily Summer (JJA) Tmin



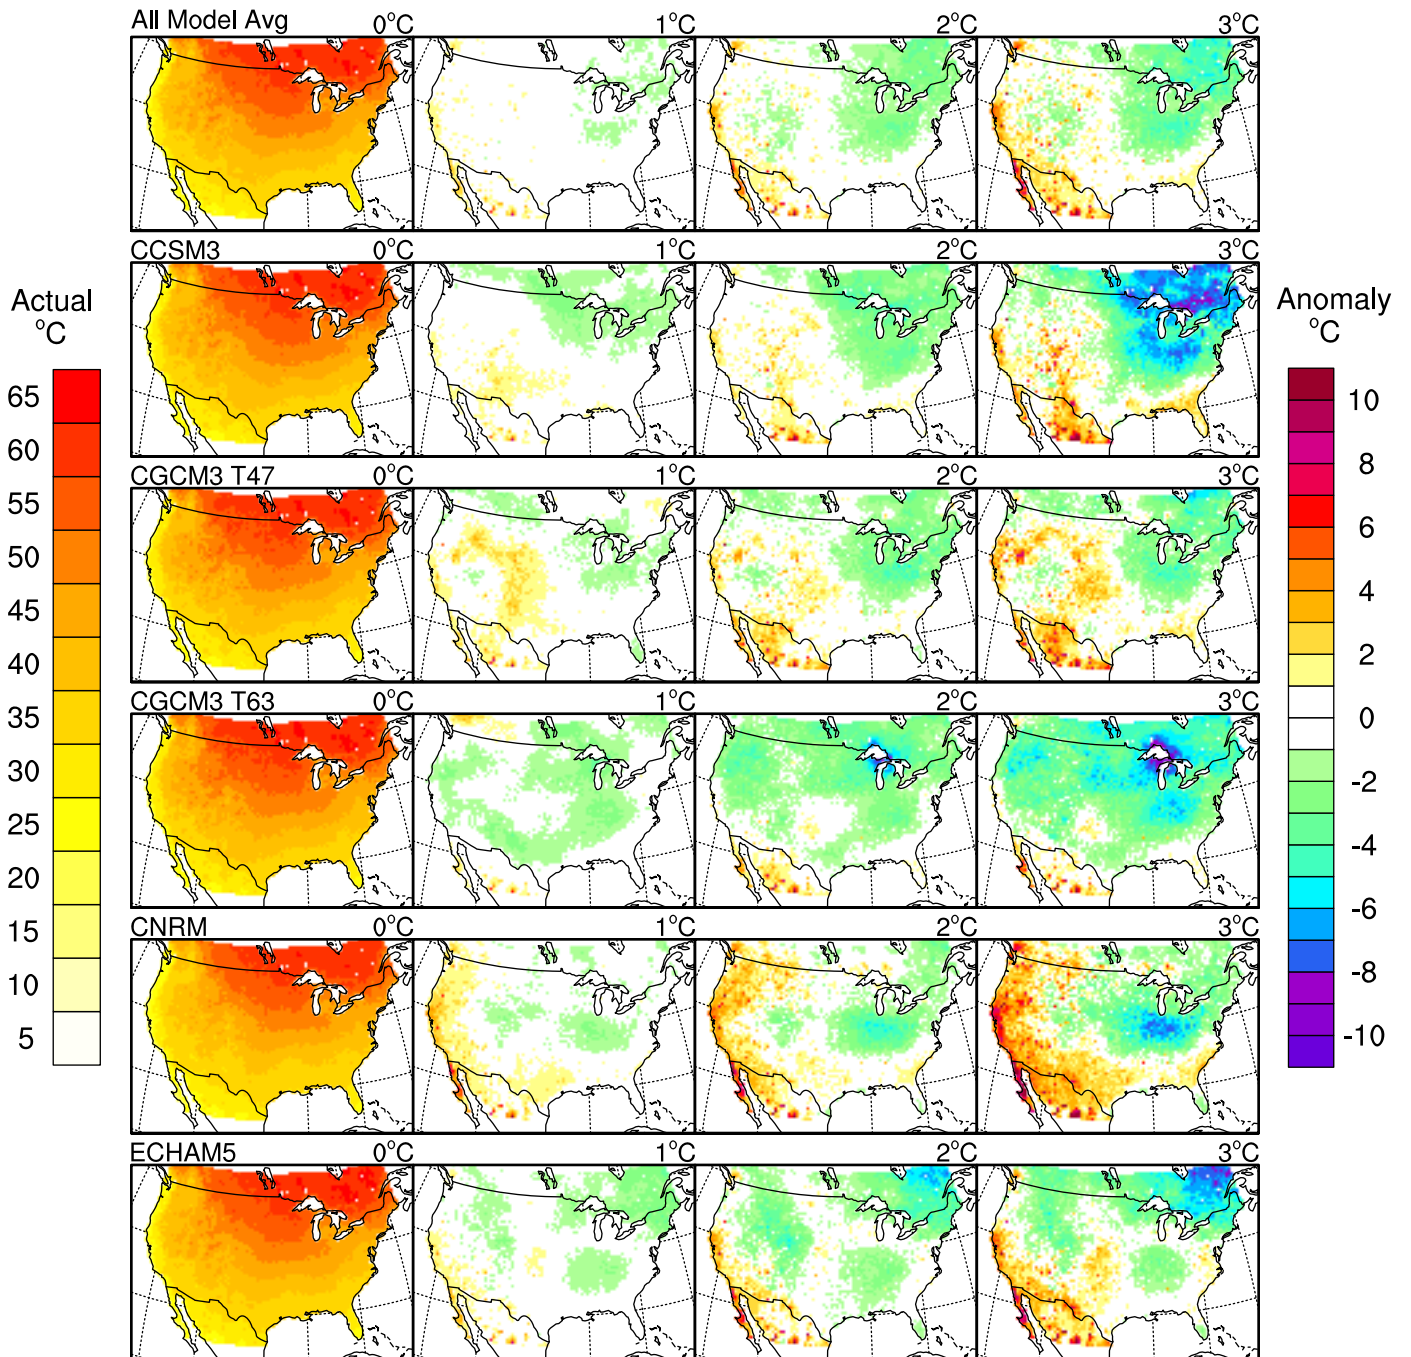
Range in Daily Fall (SON) Tmin



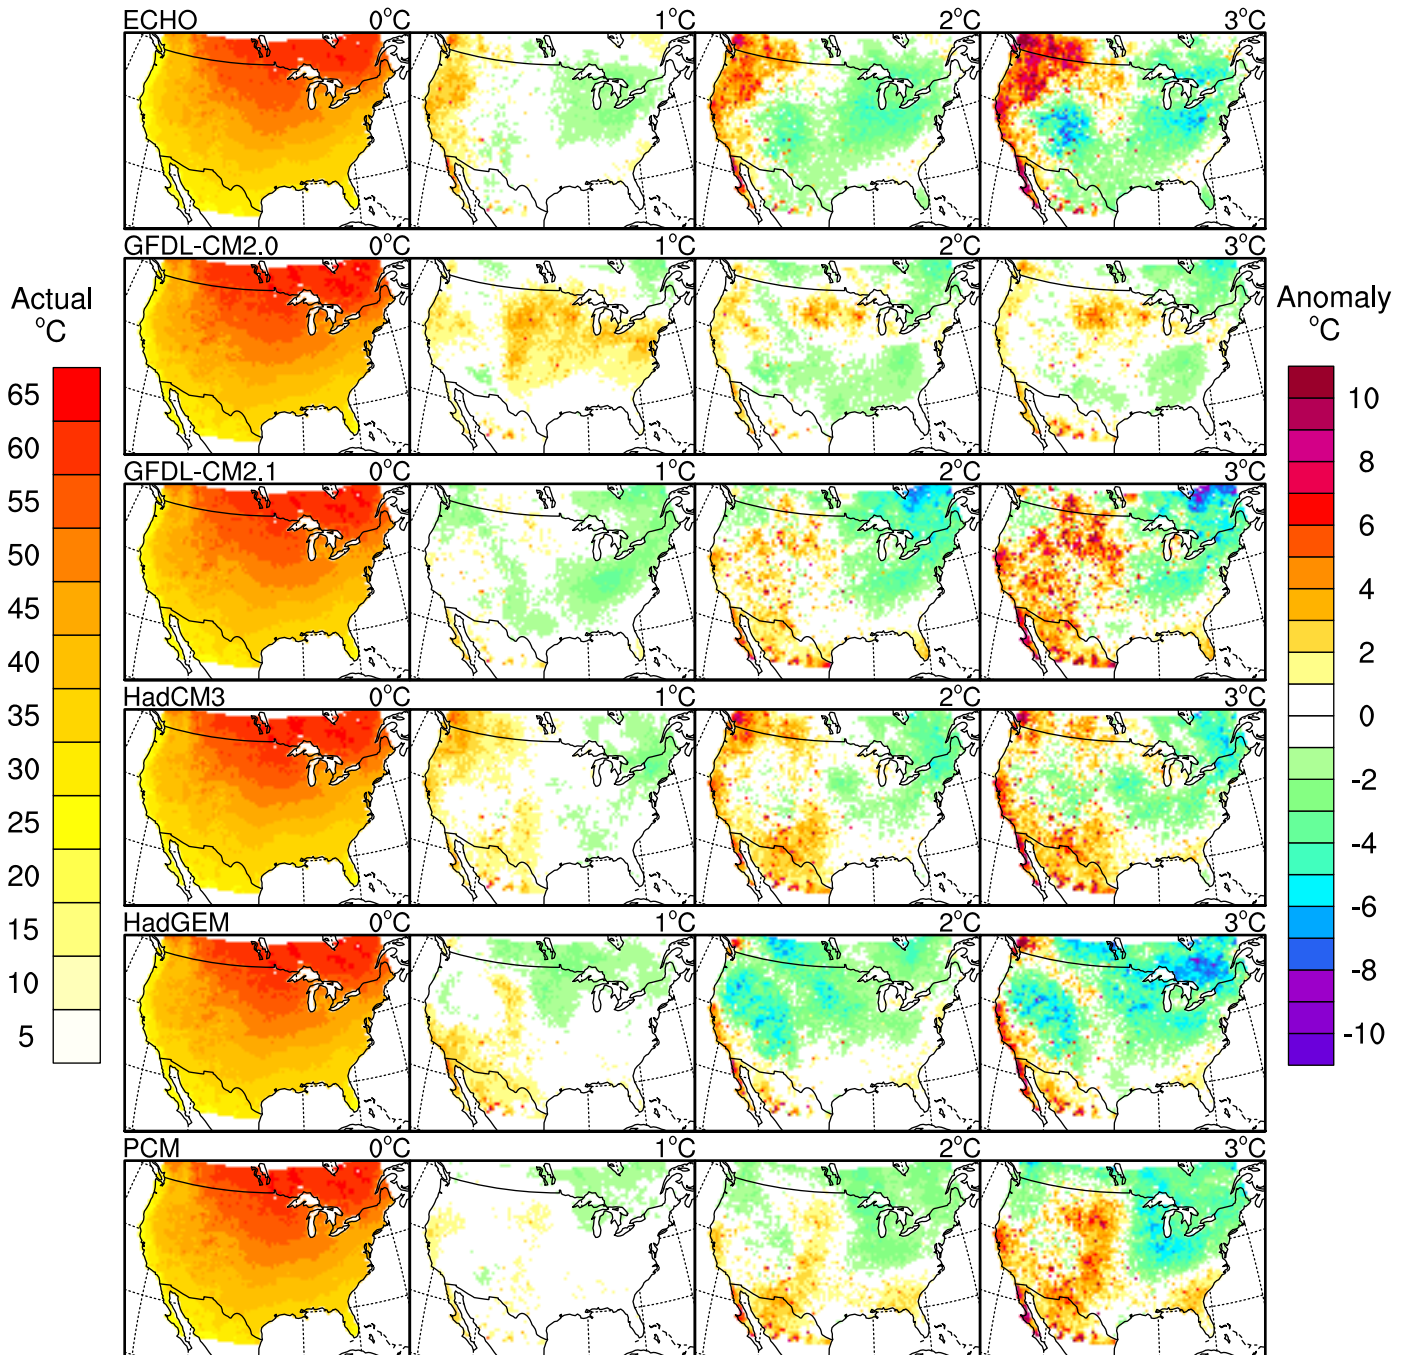
Range in Daily Fall (SON) Tmin



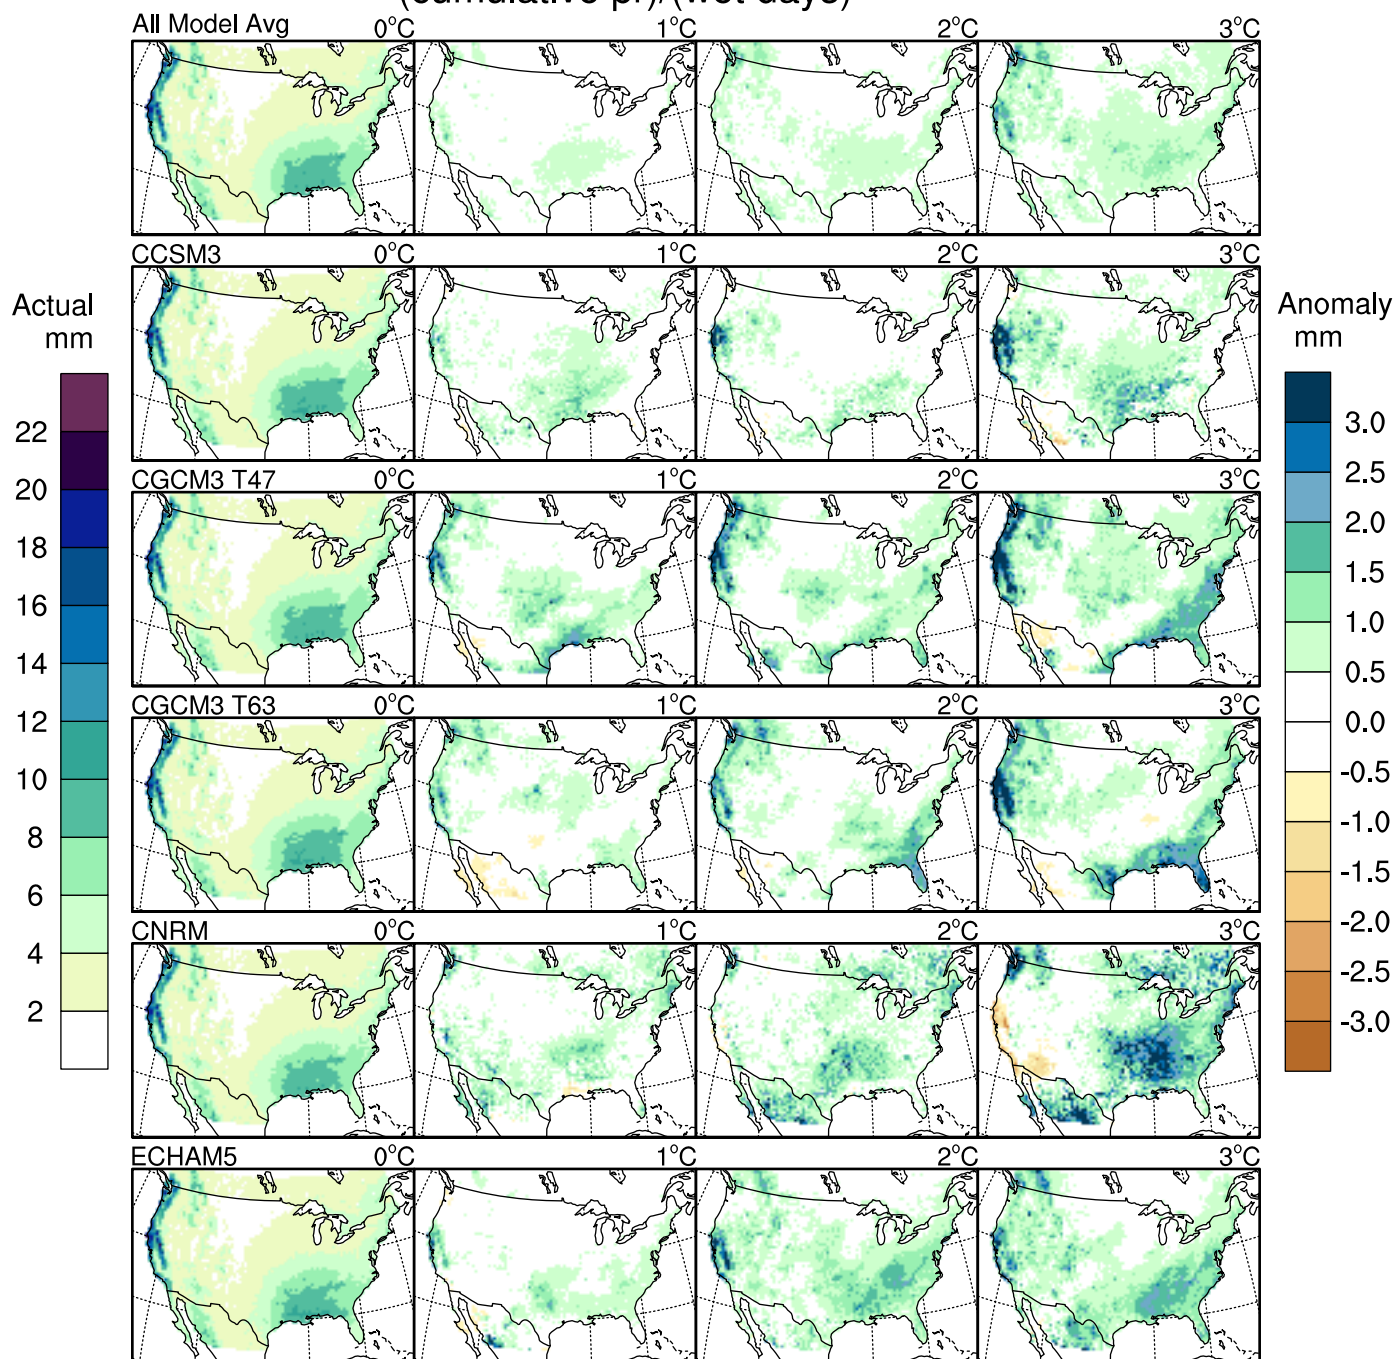
Range in Daily Annual Tmin



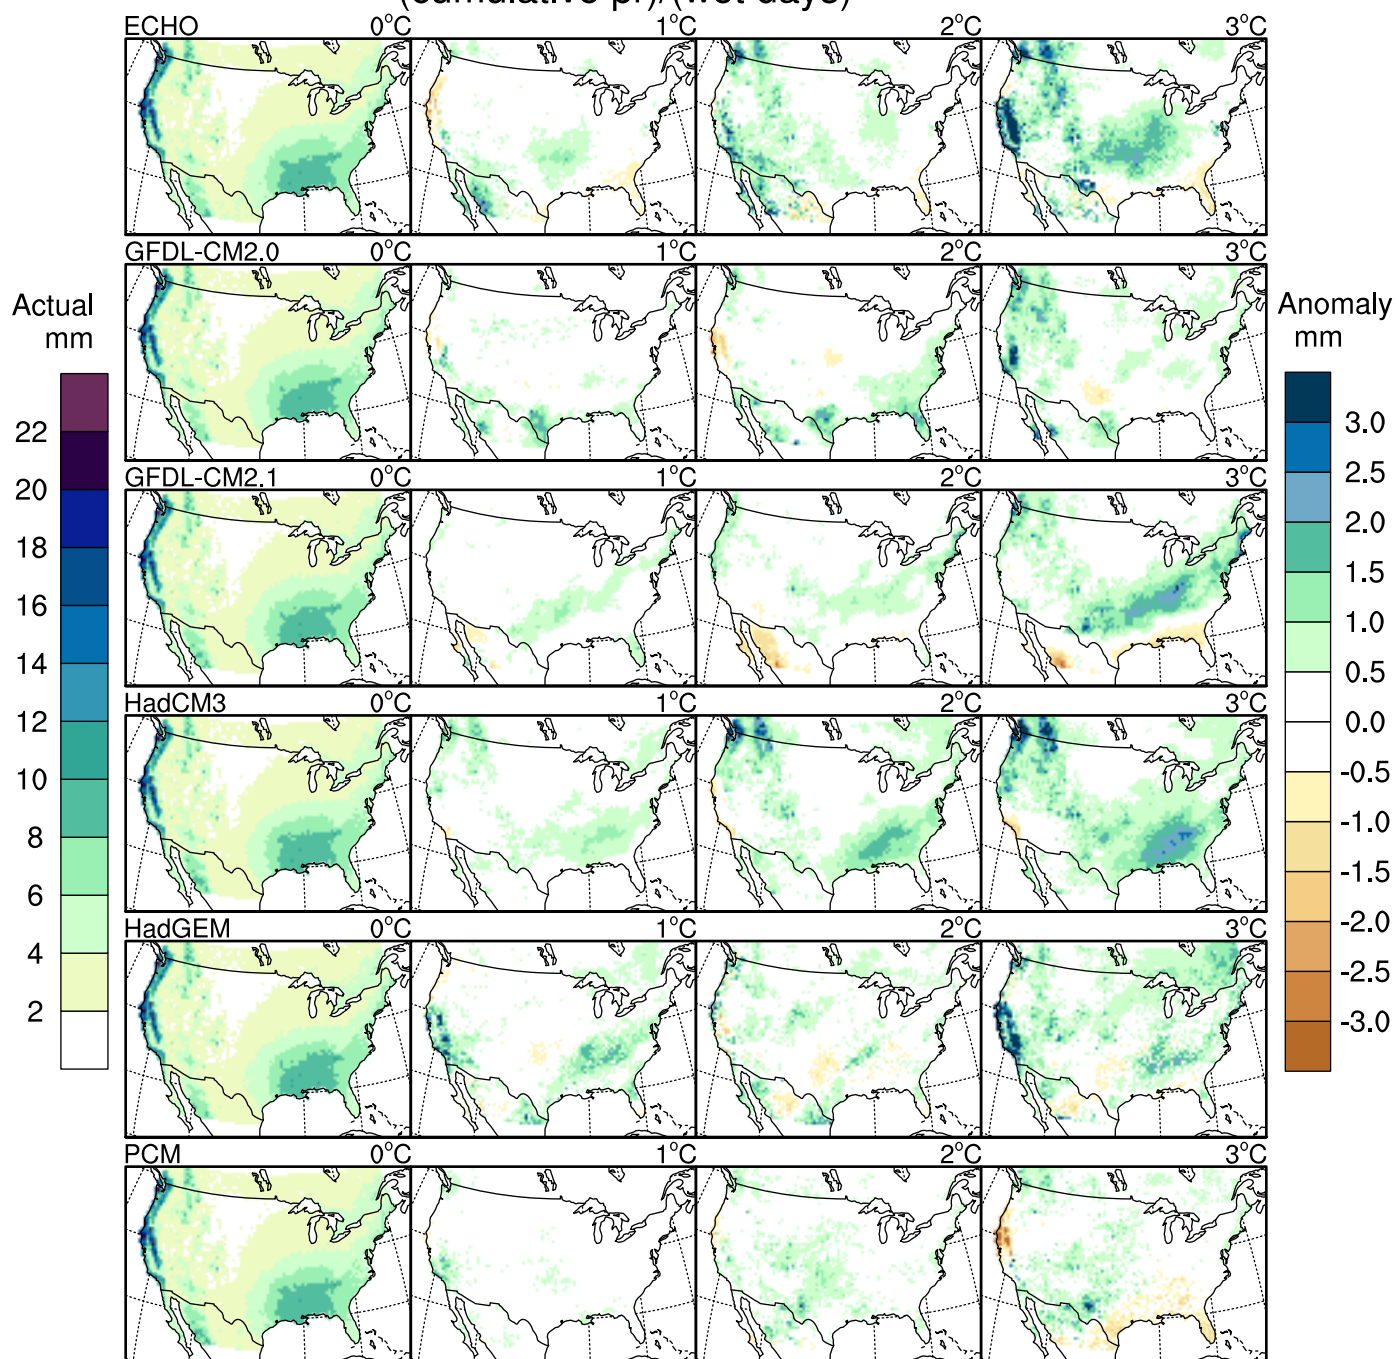
Range in Daily Annual Tmin



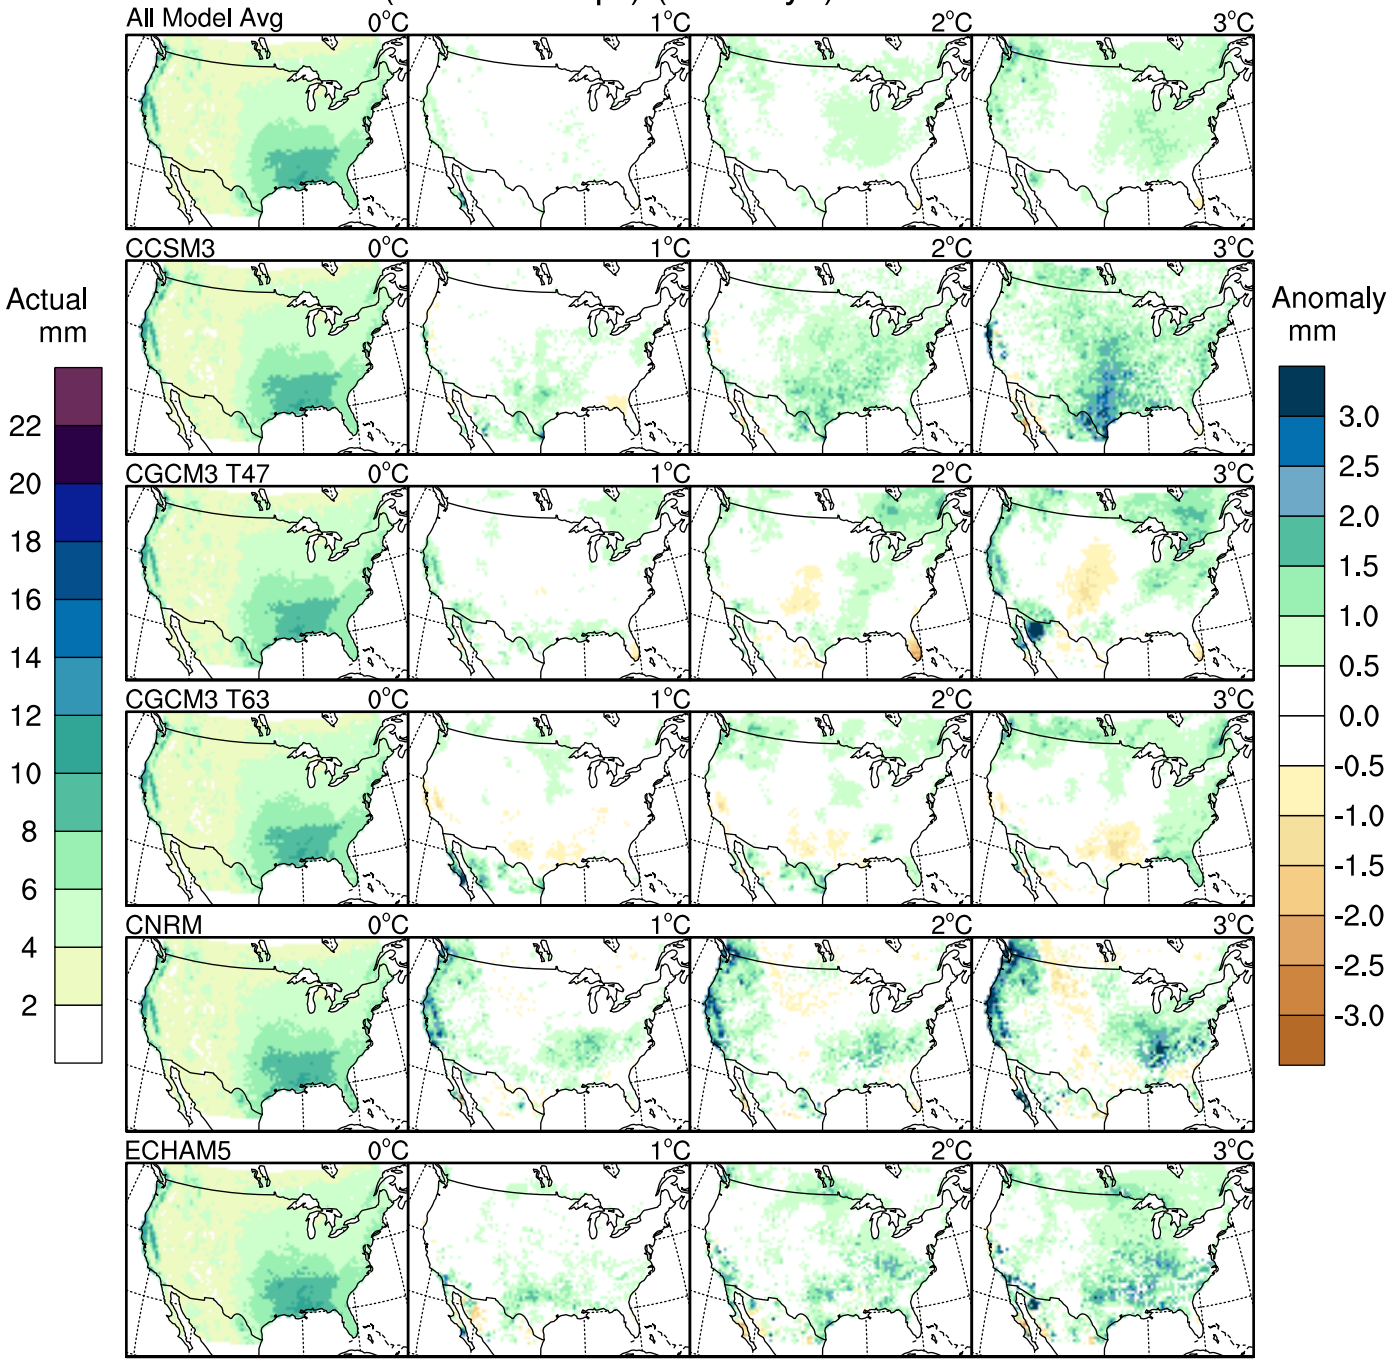
Winter (DJF) Precipitation Intensity (cumulative pr)/(wet days)



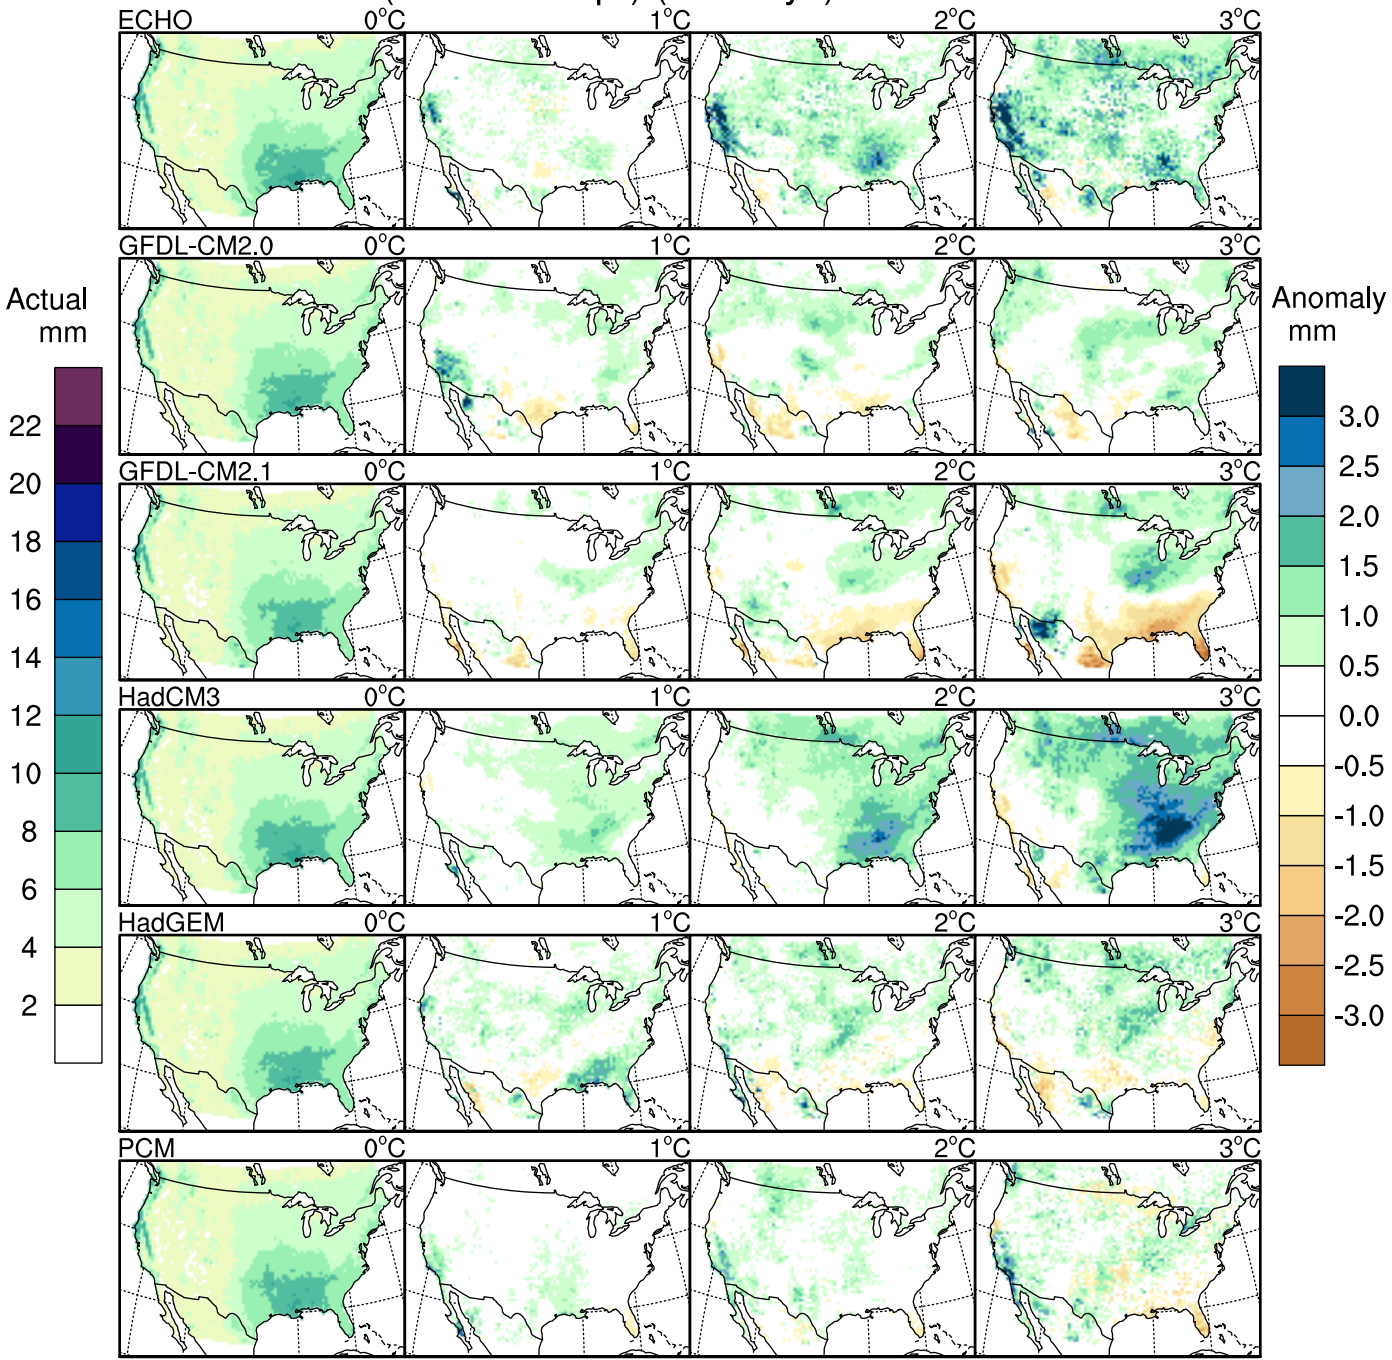
Winter (DJF) Precipitation Intensity (cumulative pr)/(wet days)



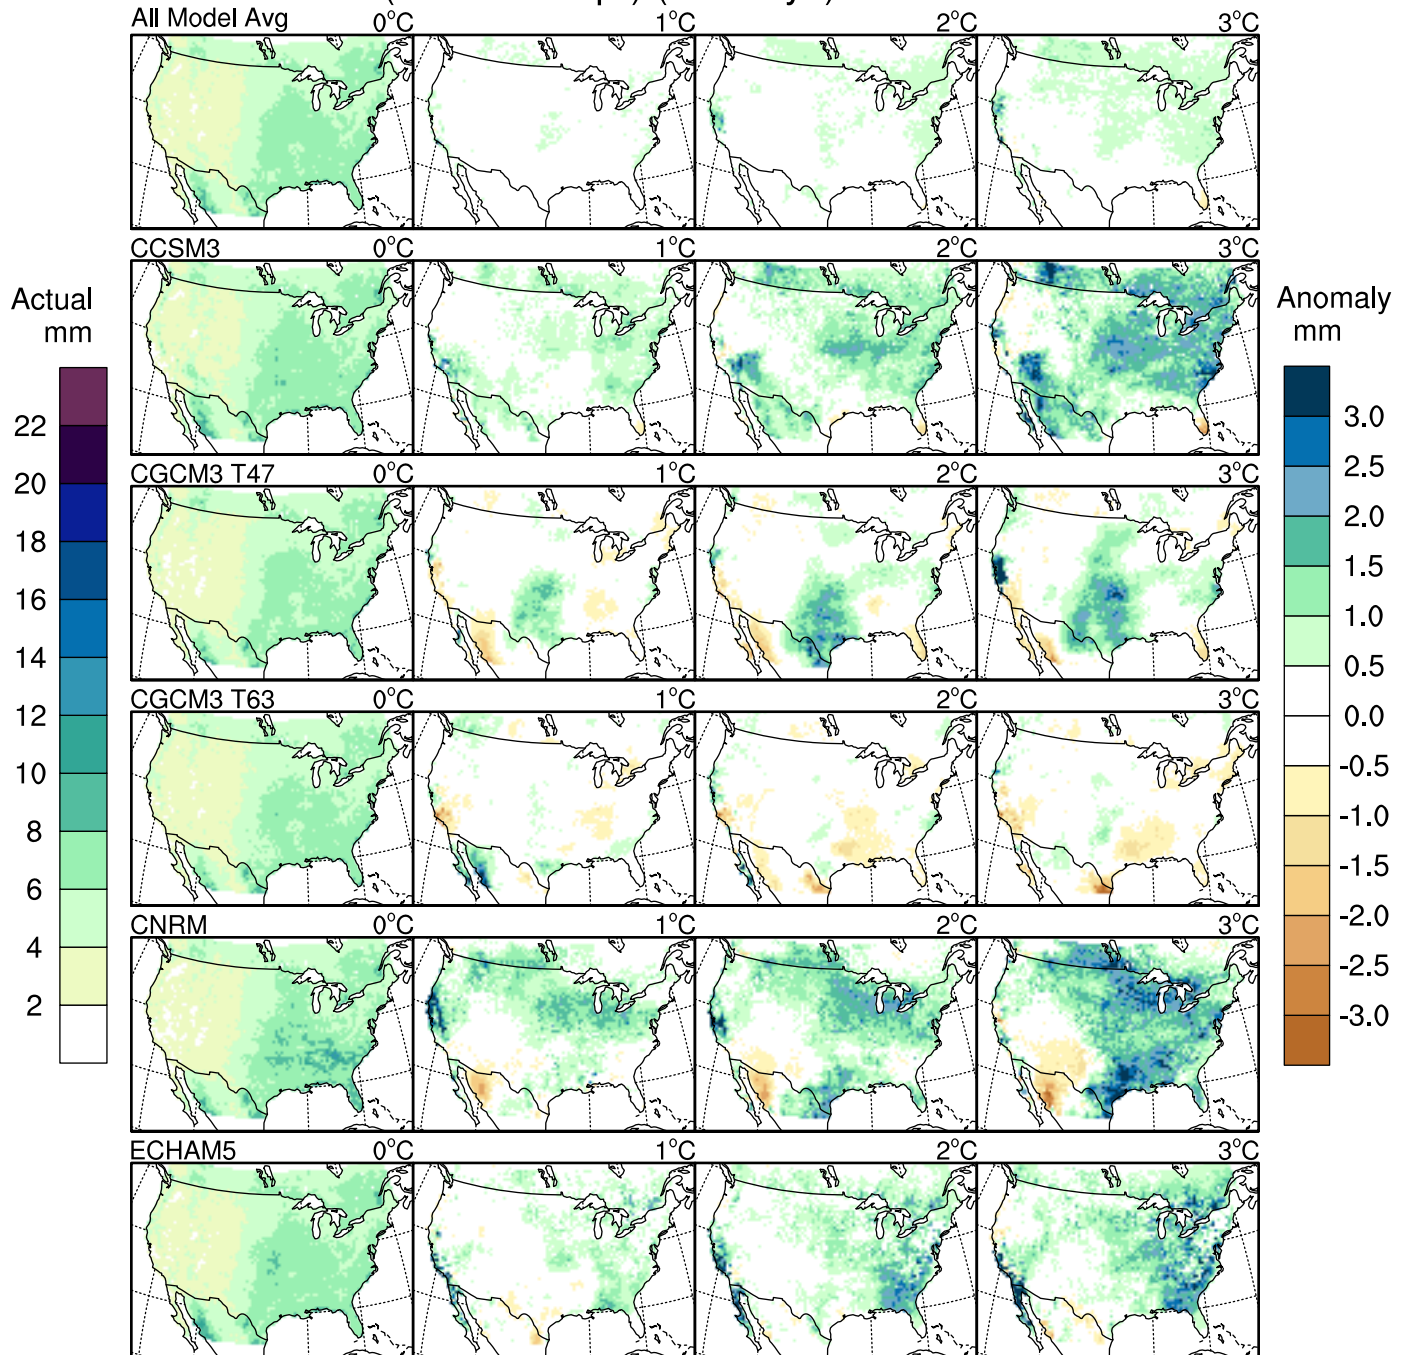
Spring (MAM) Precipitation Intensity
(cumulative pr)/(wet days)



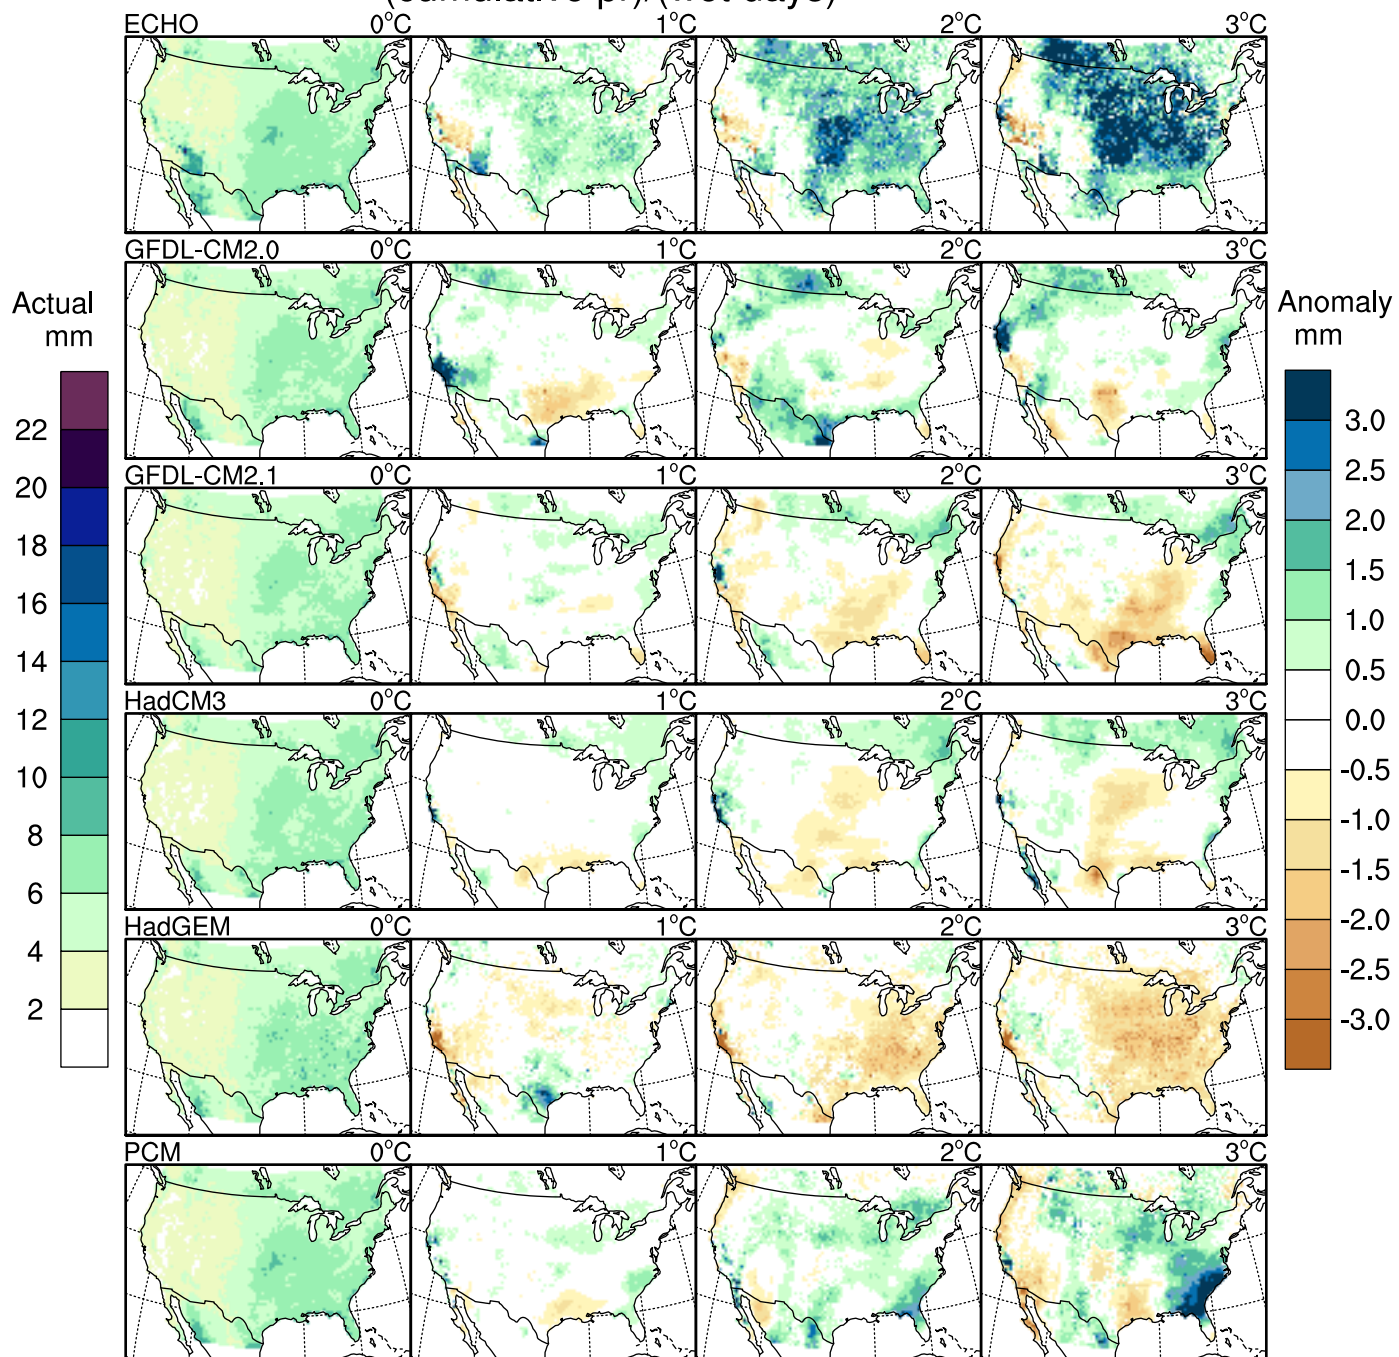
Spring (MAM) Precipitation Intensity
(cumulative pr)/(wet days)



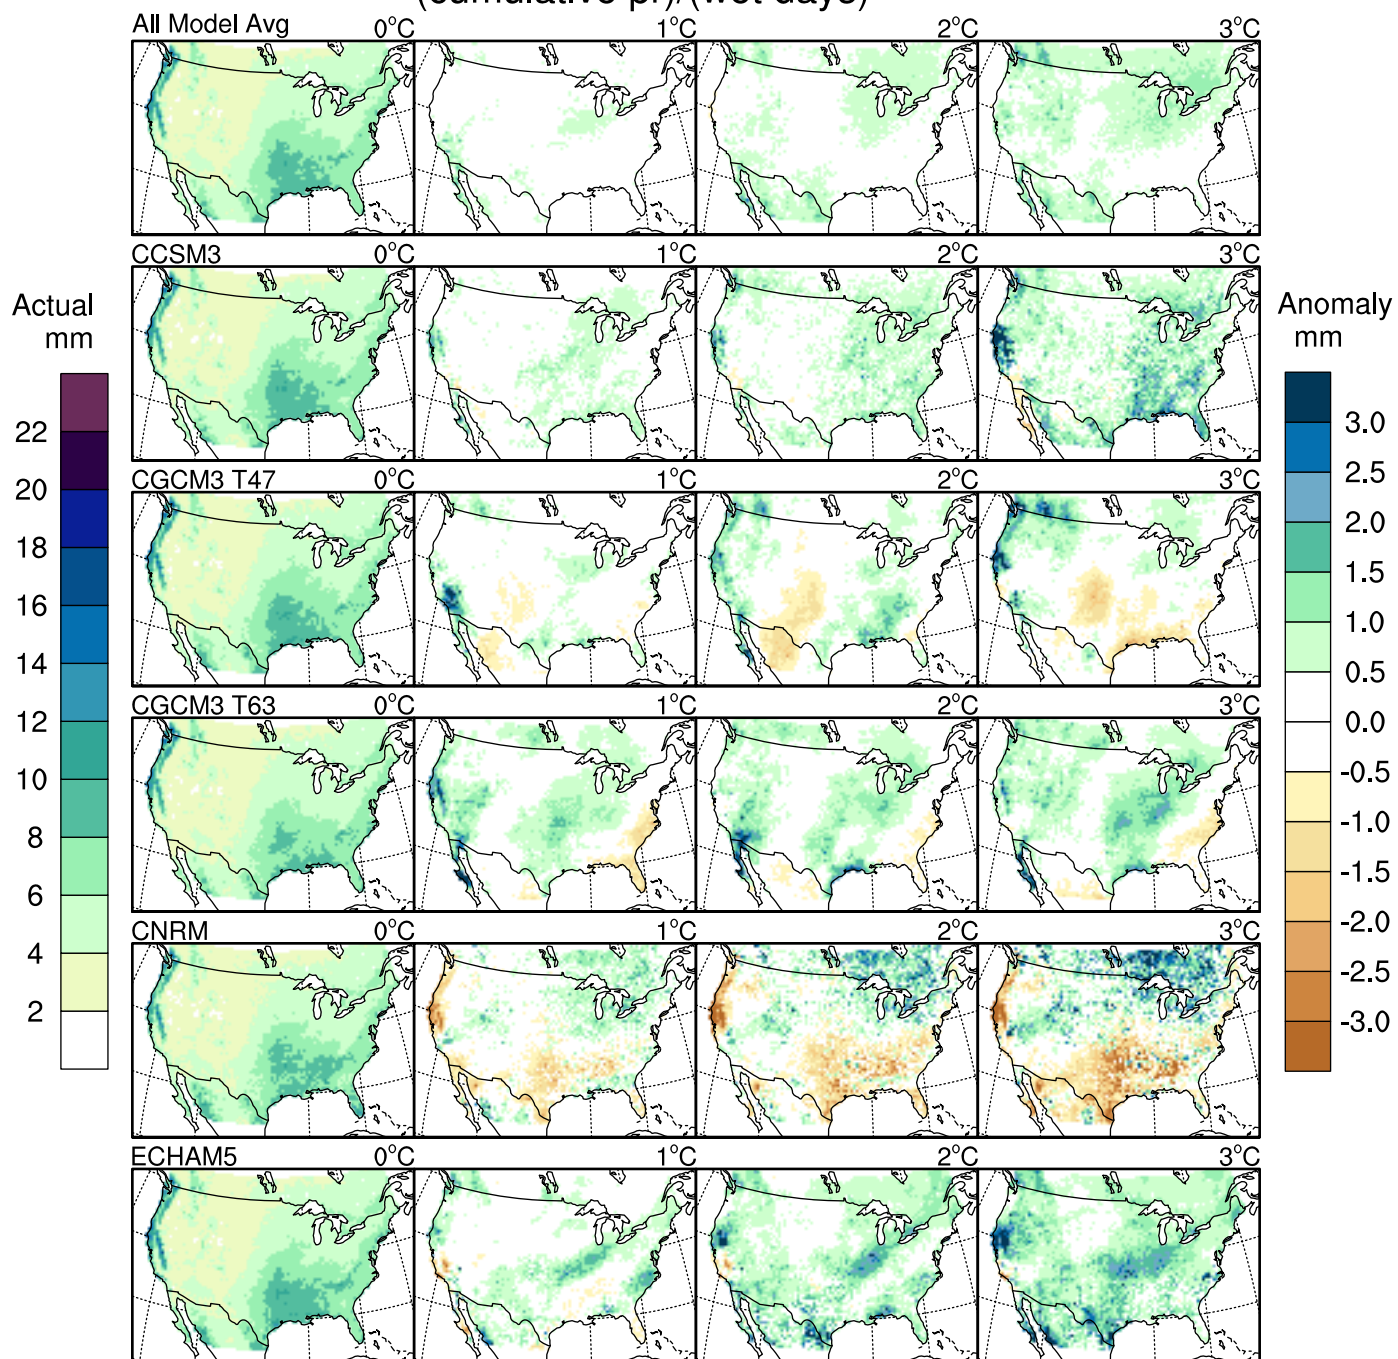
Summer (JJA) Precipitation Intensity (cumulative pr)/(wet days)



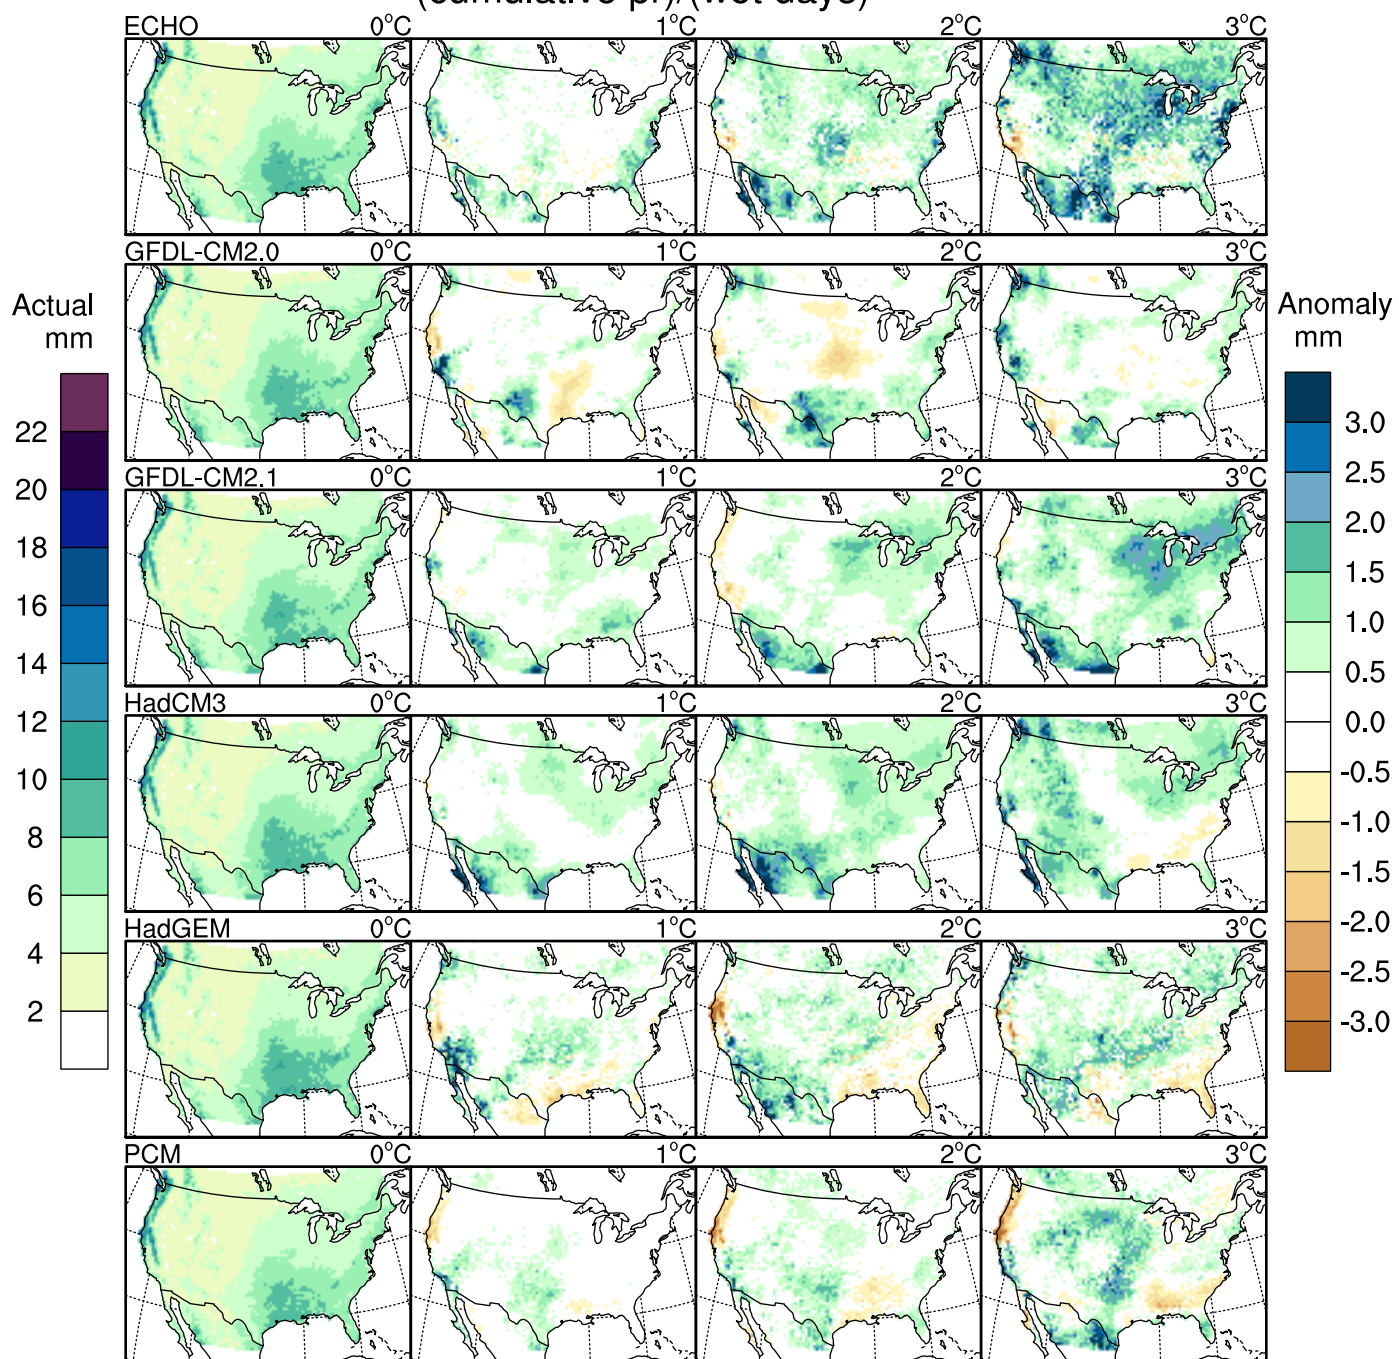
Summer (JJA) Precipitation Intensity (cumulative pr)/(wet days)



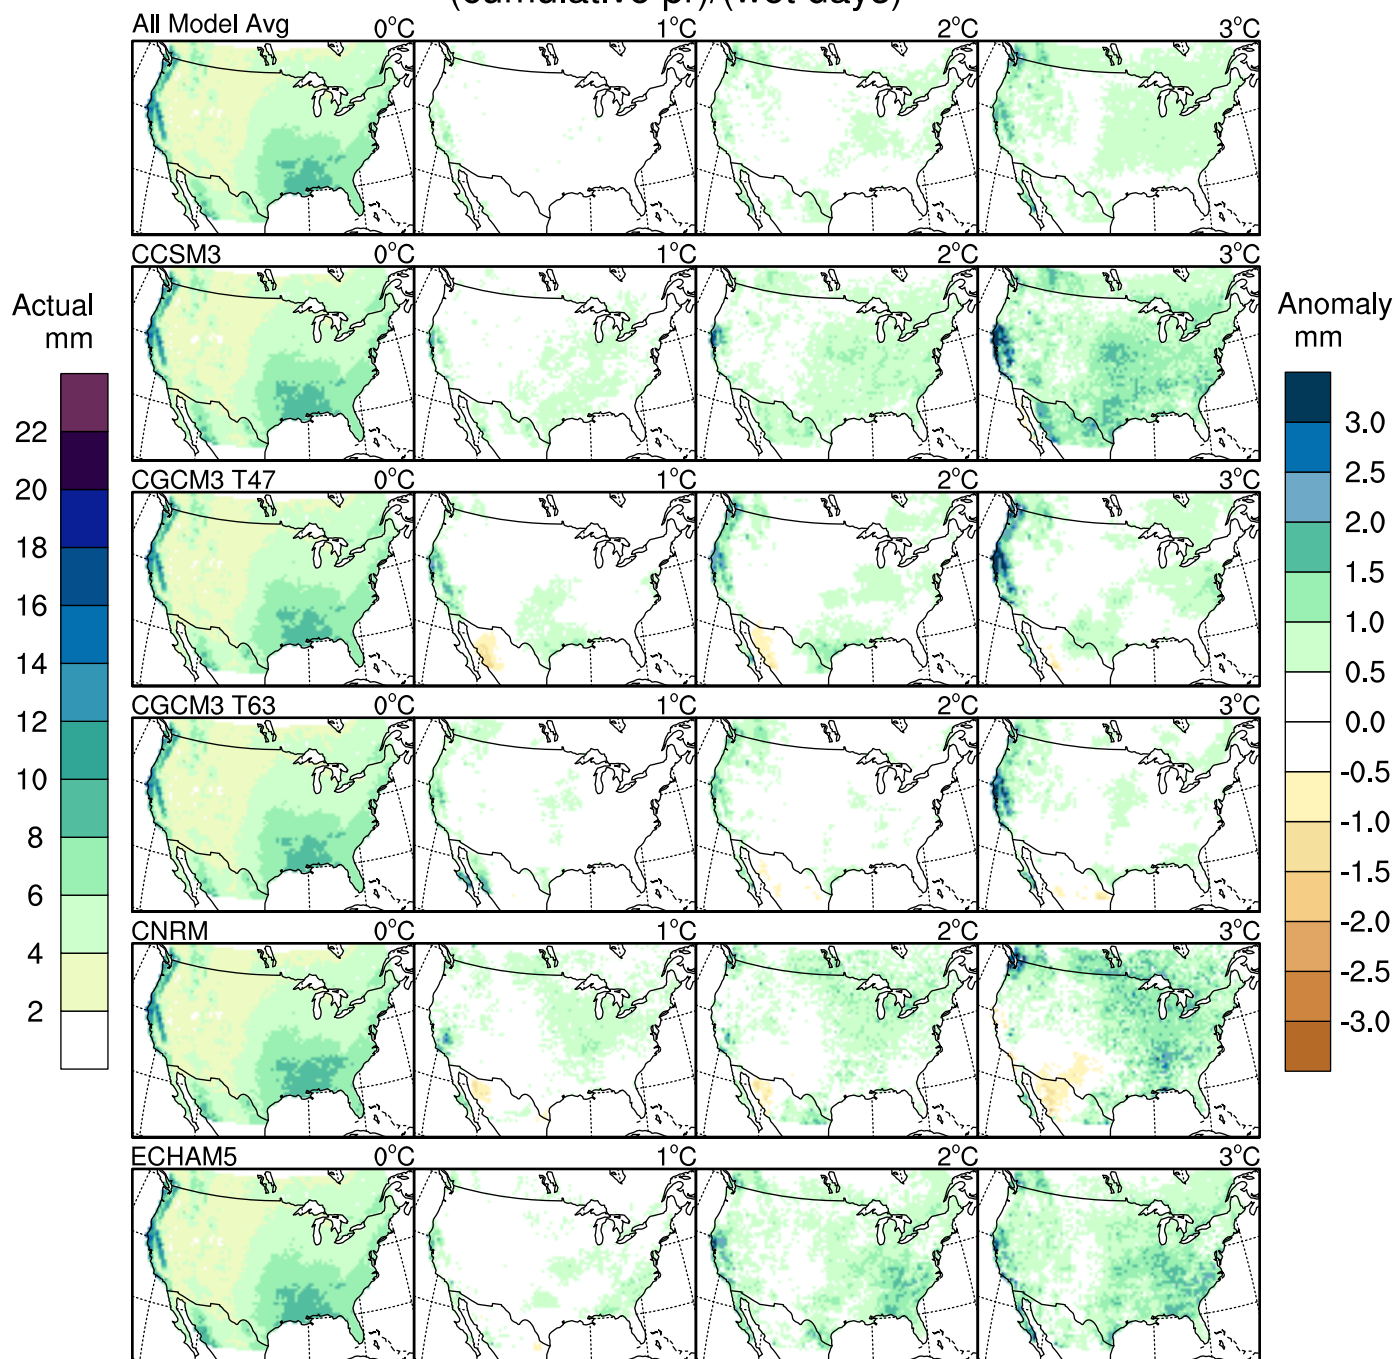
Fall (SON) Precipitation Intensity (cumulative pr)/(wet days)



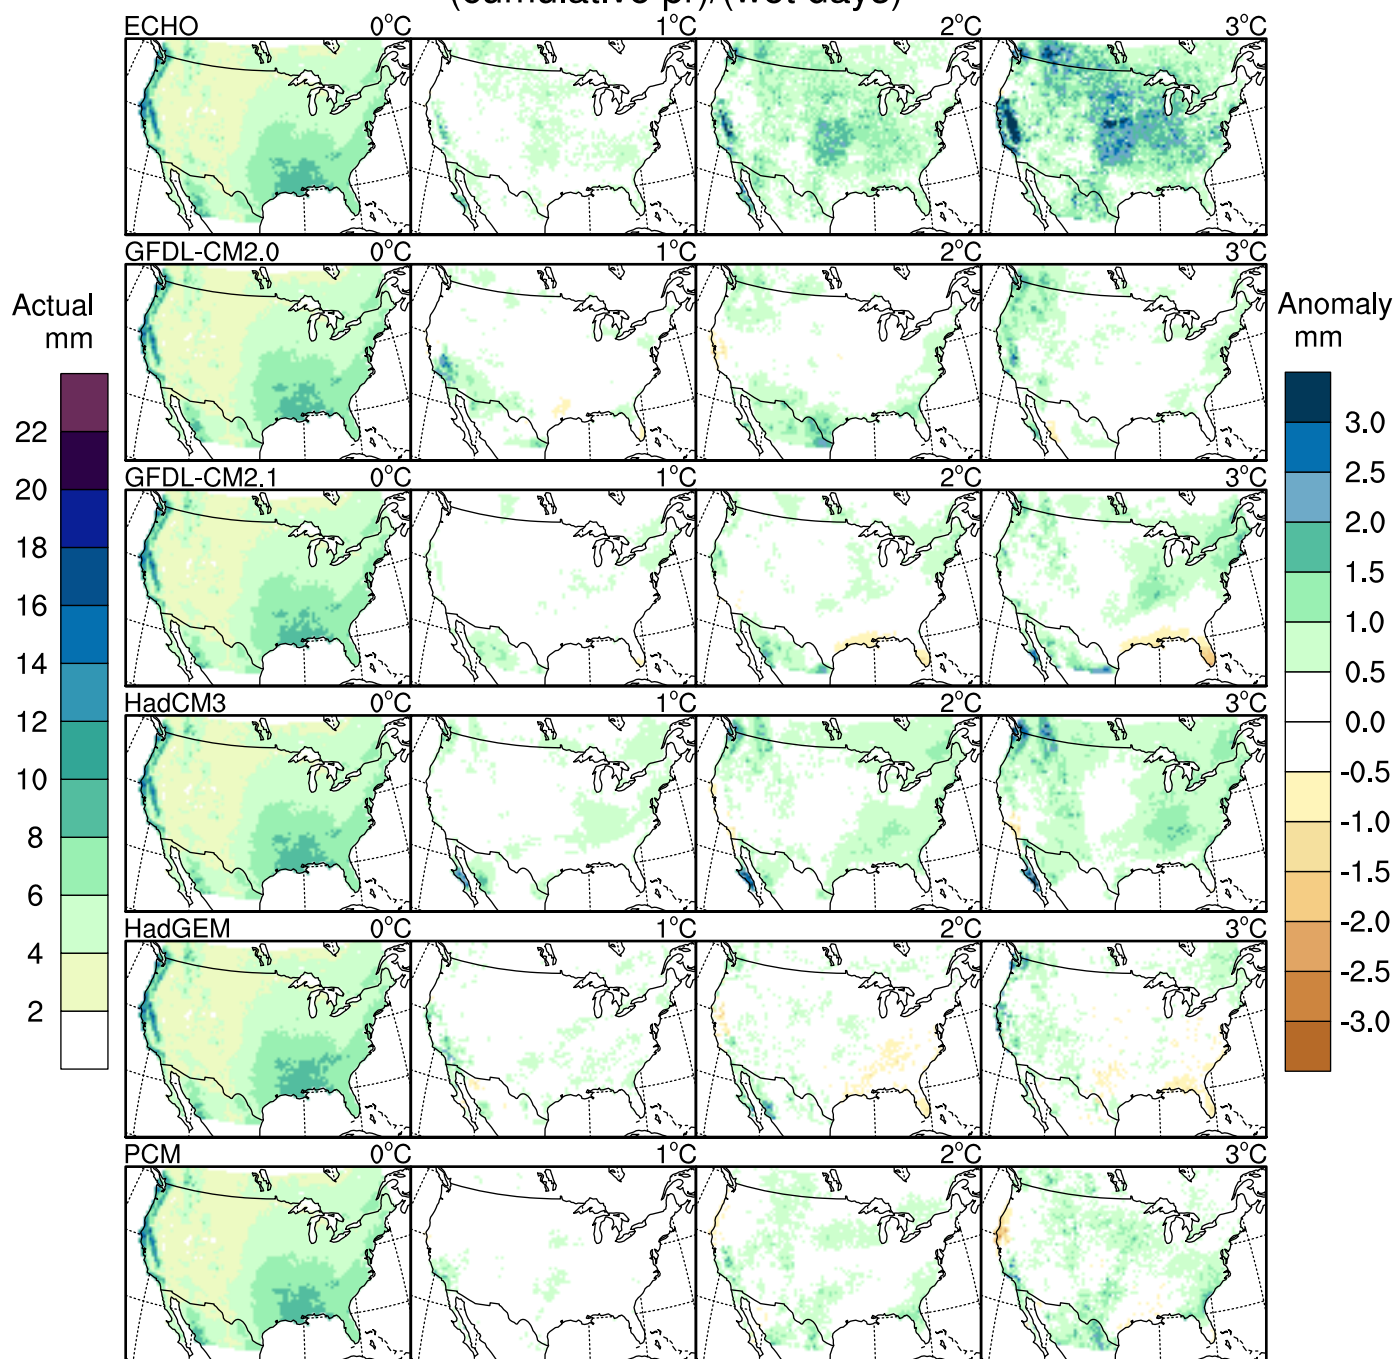
Fall (SON) Precipitation Intensity (cumulative pr)/(wet days)



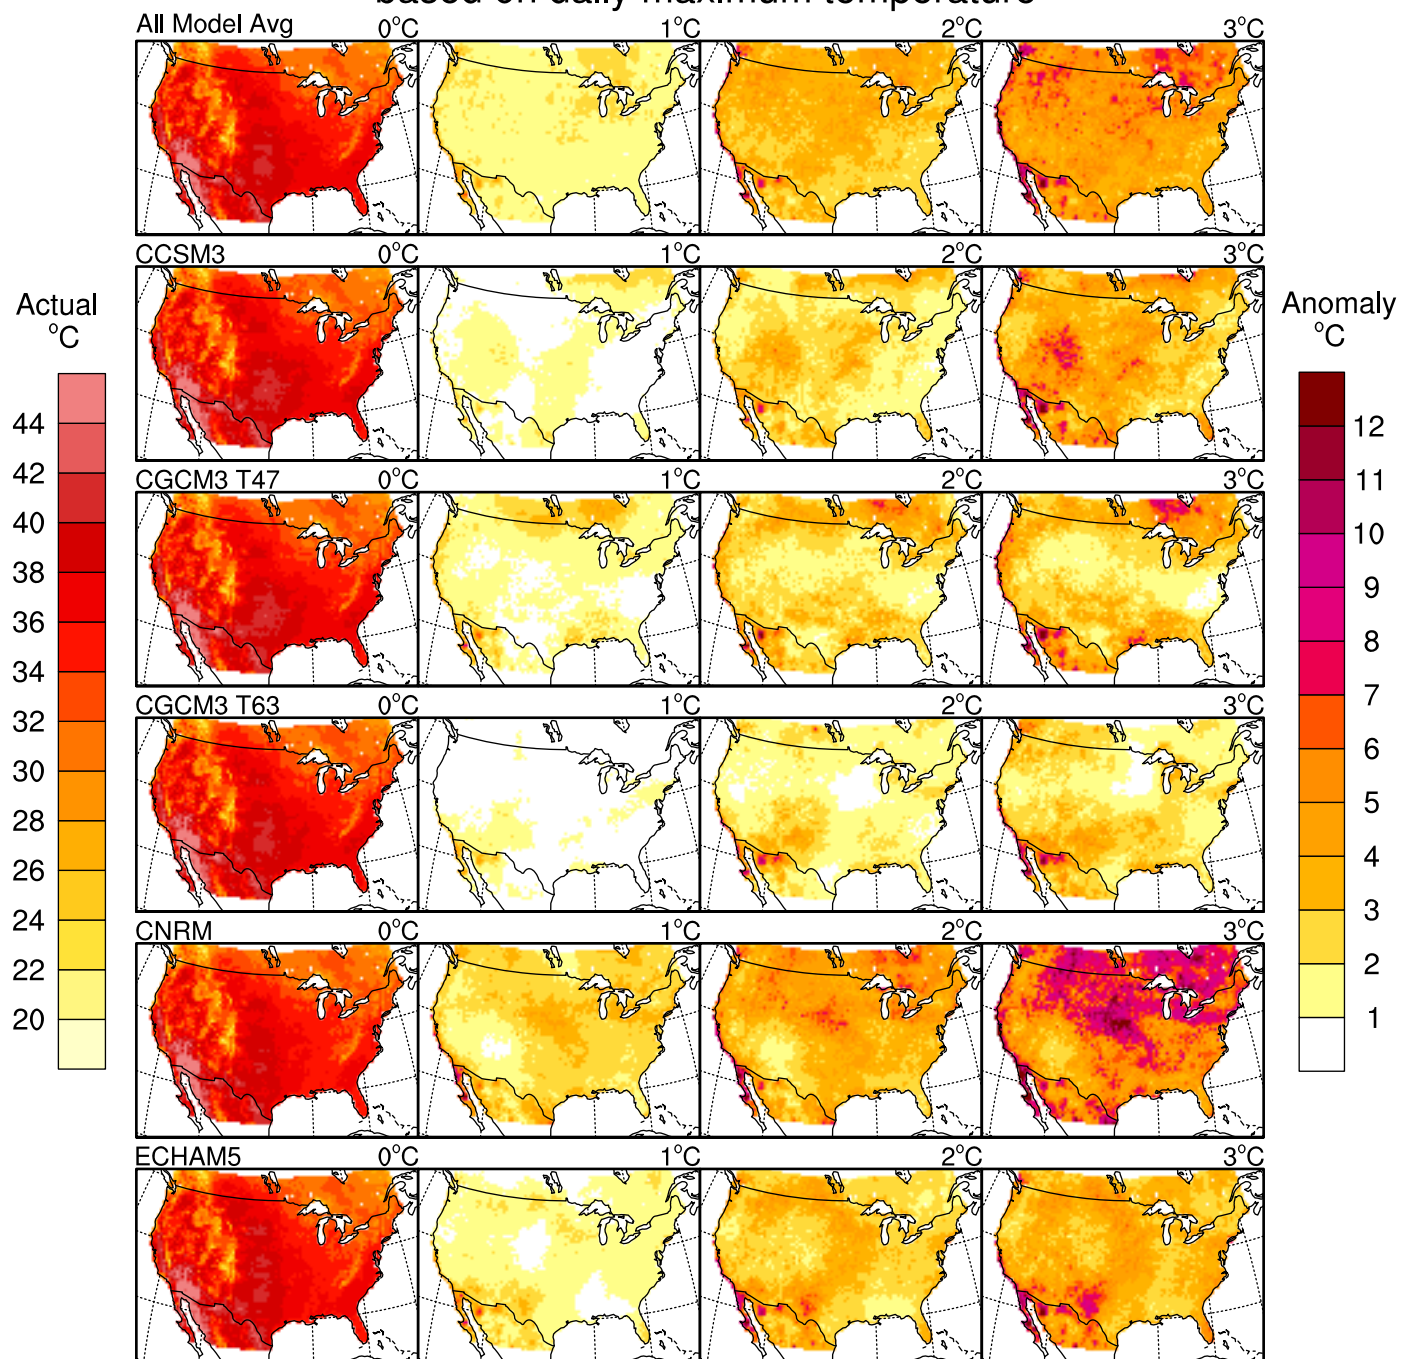
Annual Precipitation Intensity (cumulative pr)/(wet days)



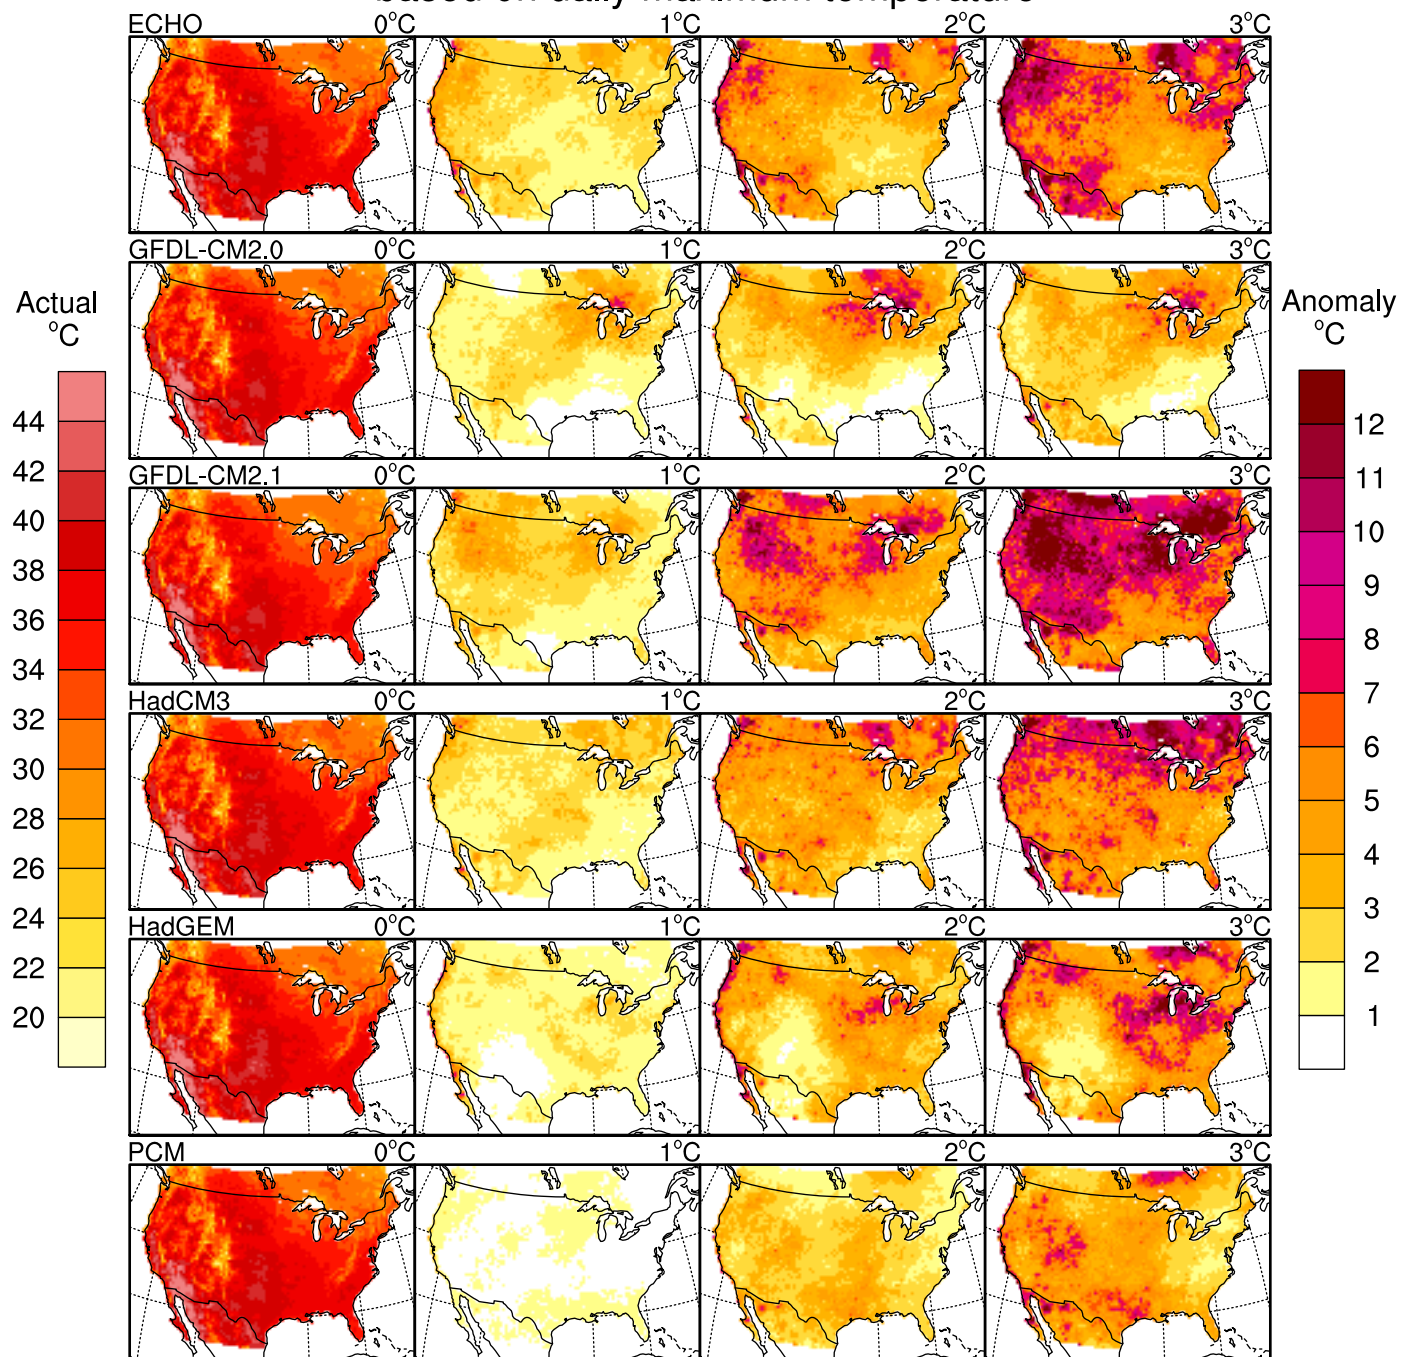
Annual Precipitation Intensity (cumulative pr)/(wet days)



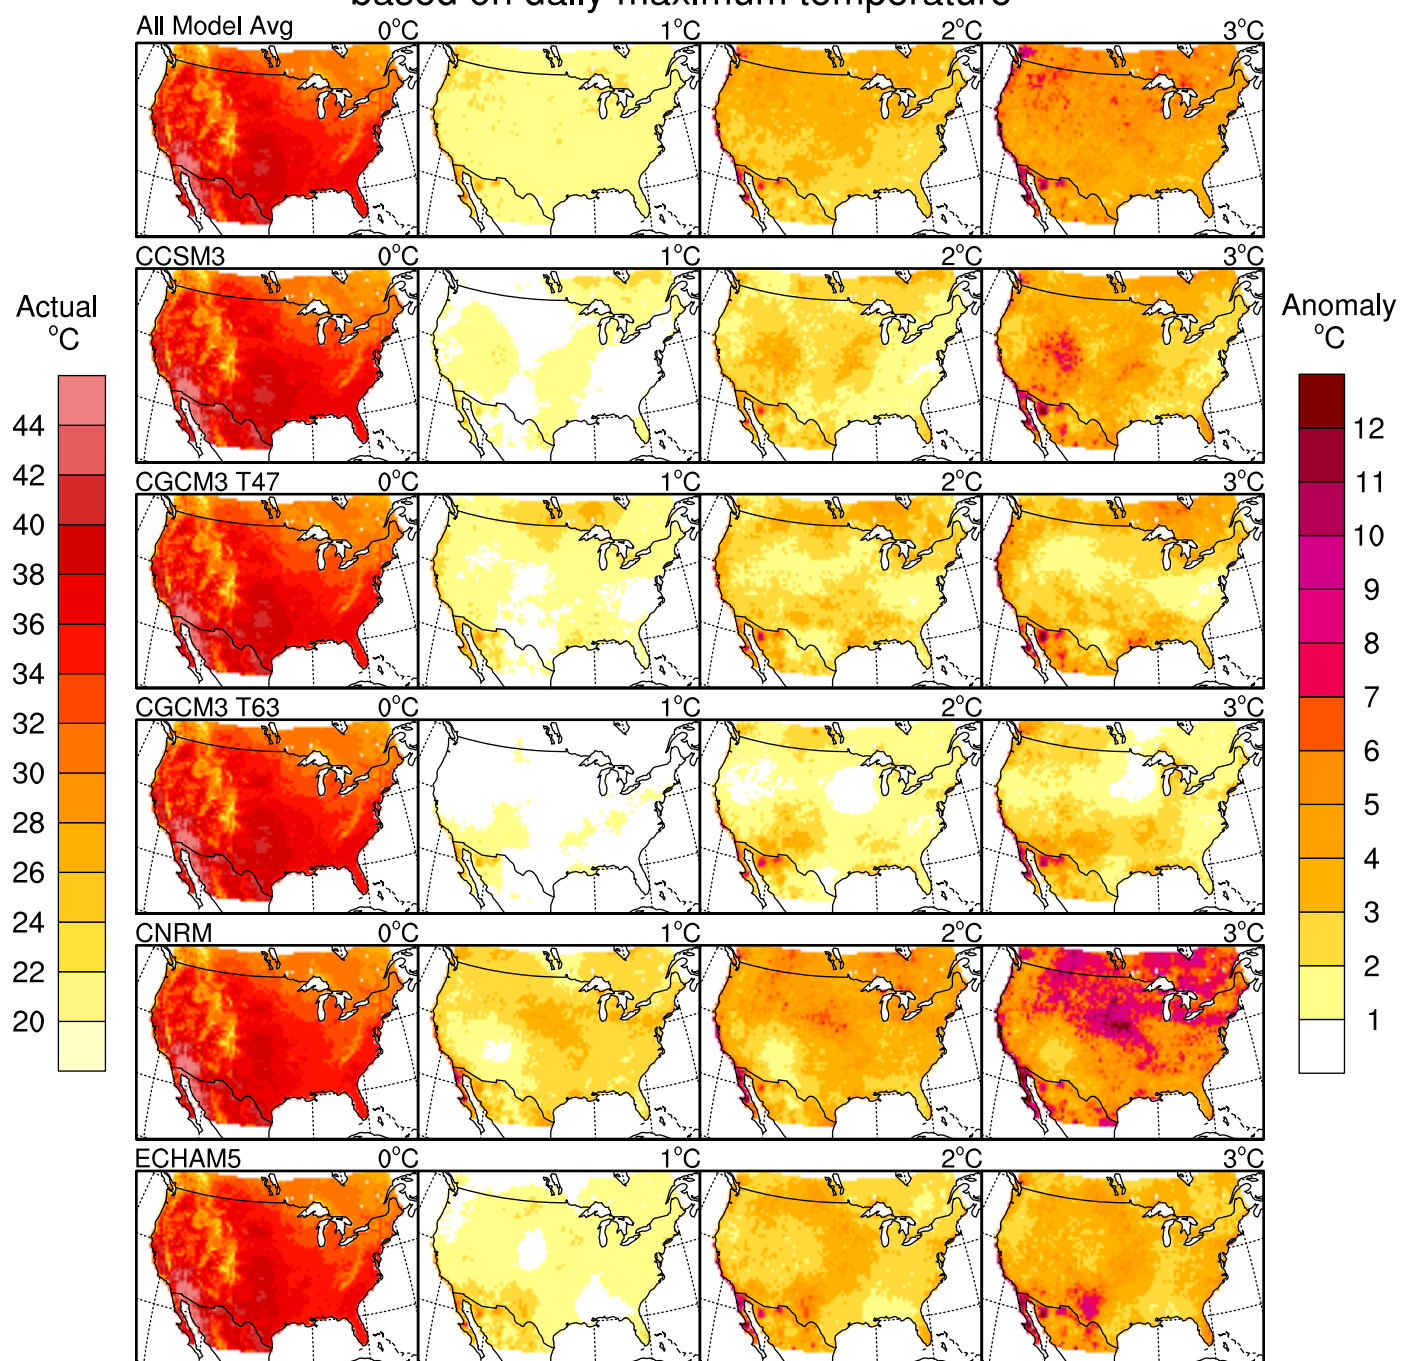
Temperature of the Hottest Day based on daily maximum temperature



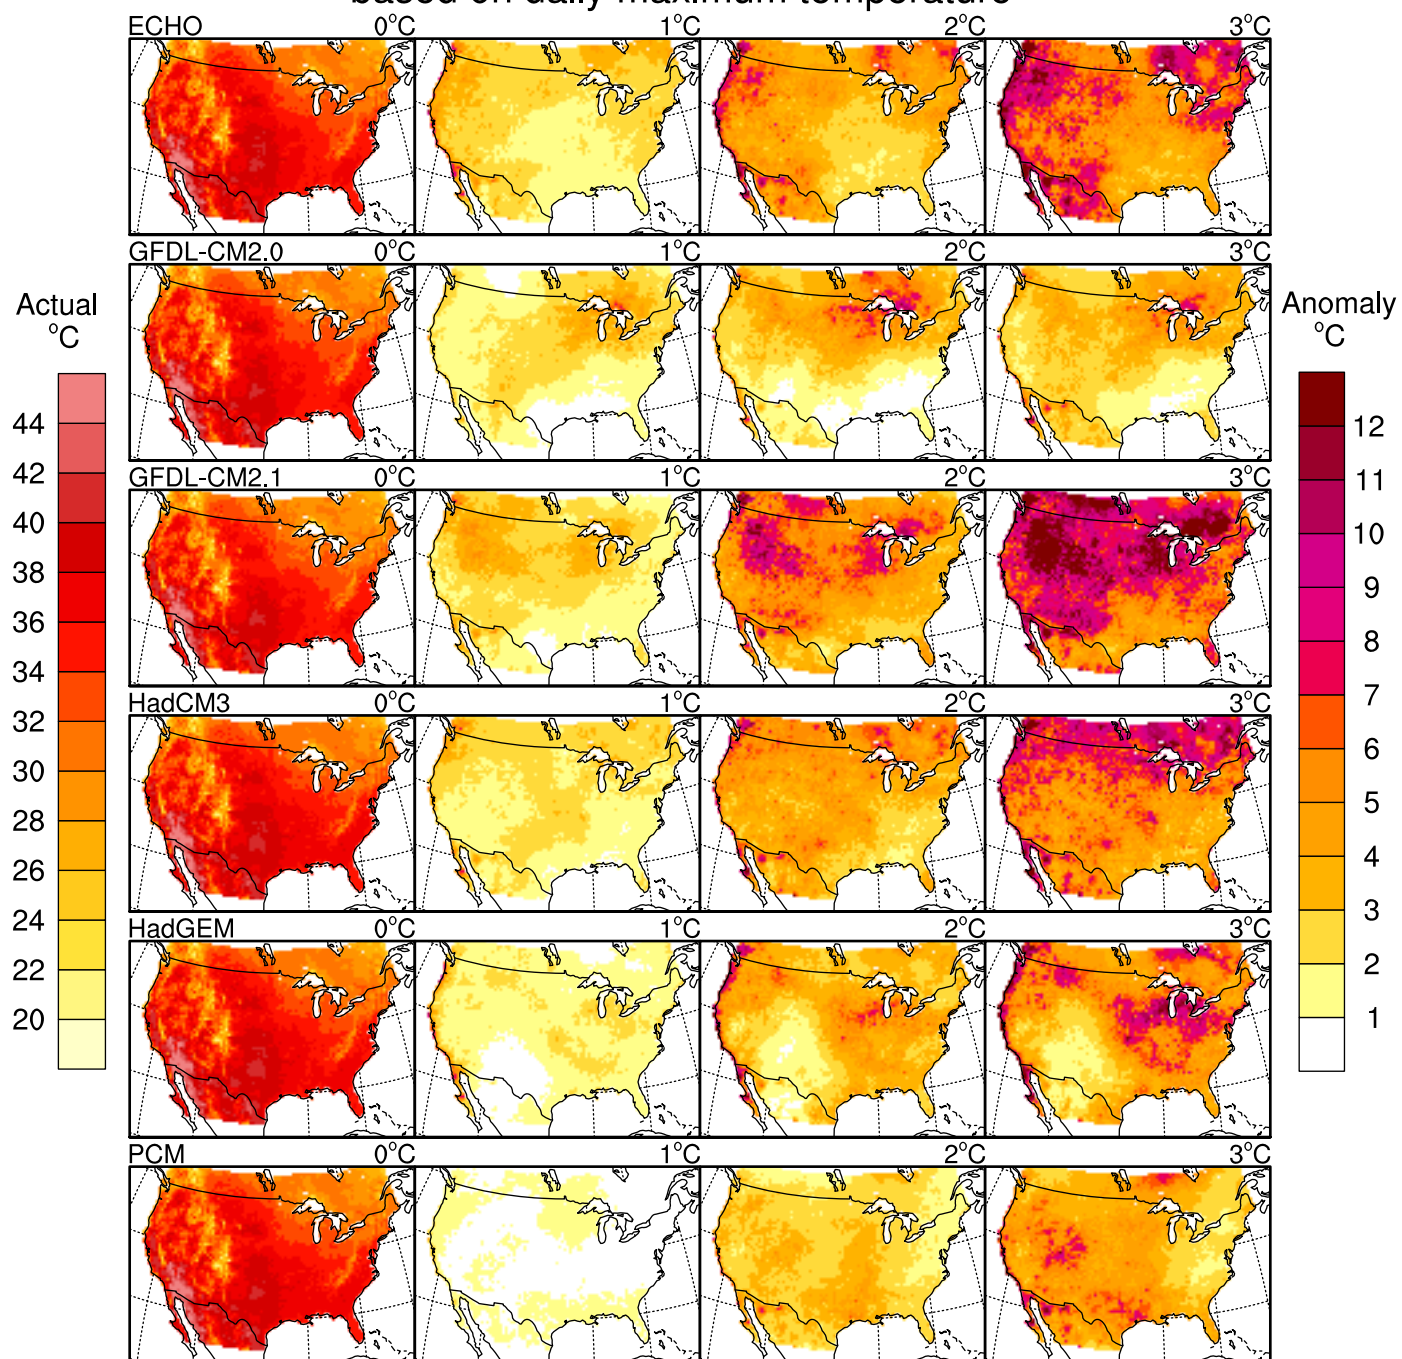
Temperature of the Hottest Day based on daily maximum temperature



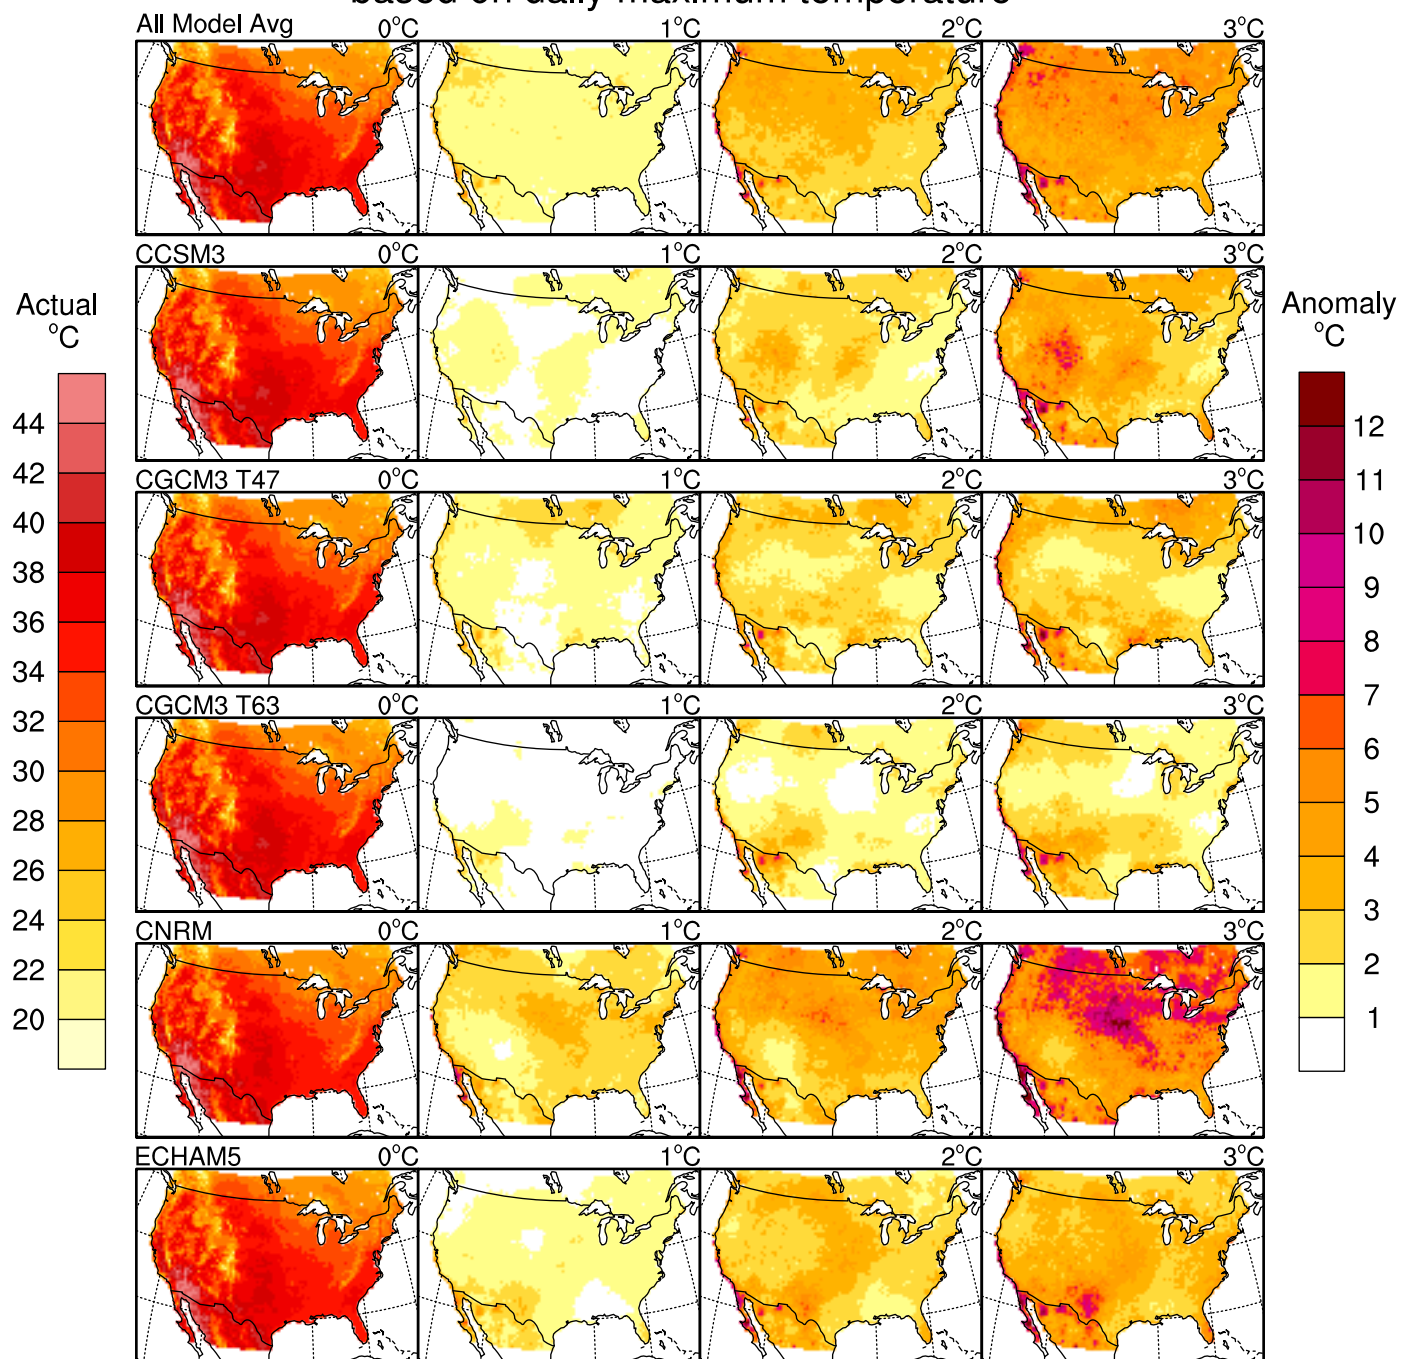
Temperature of the Hottest 3-Day Period based on daily maximum temperature



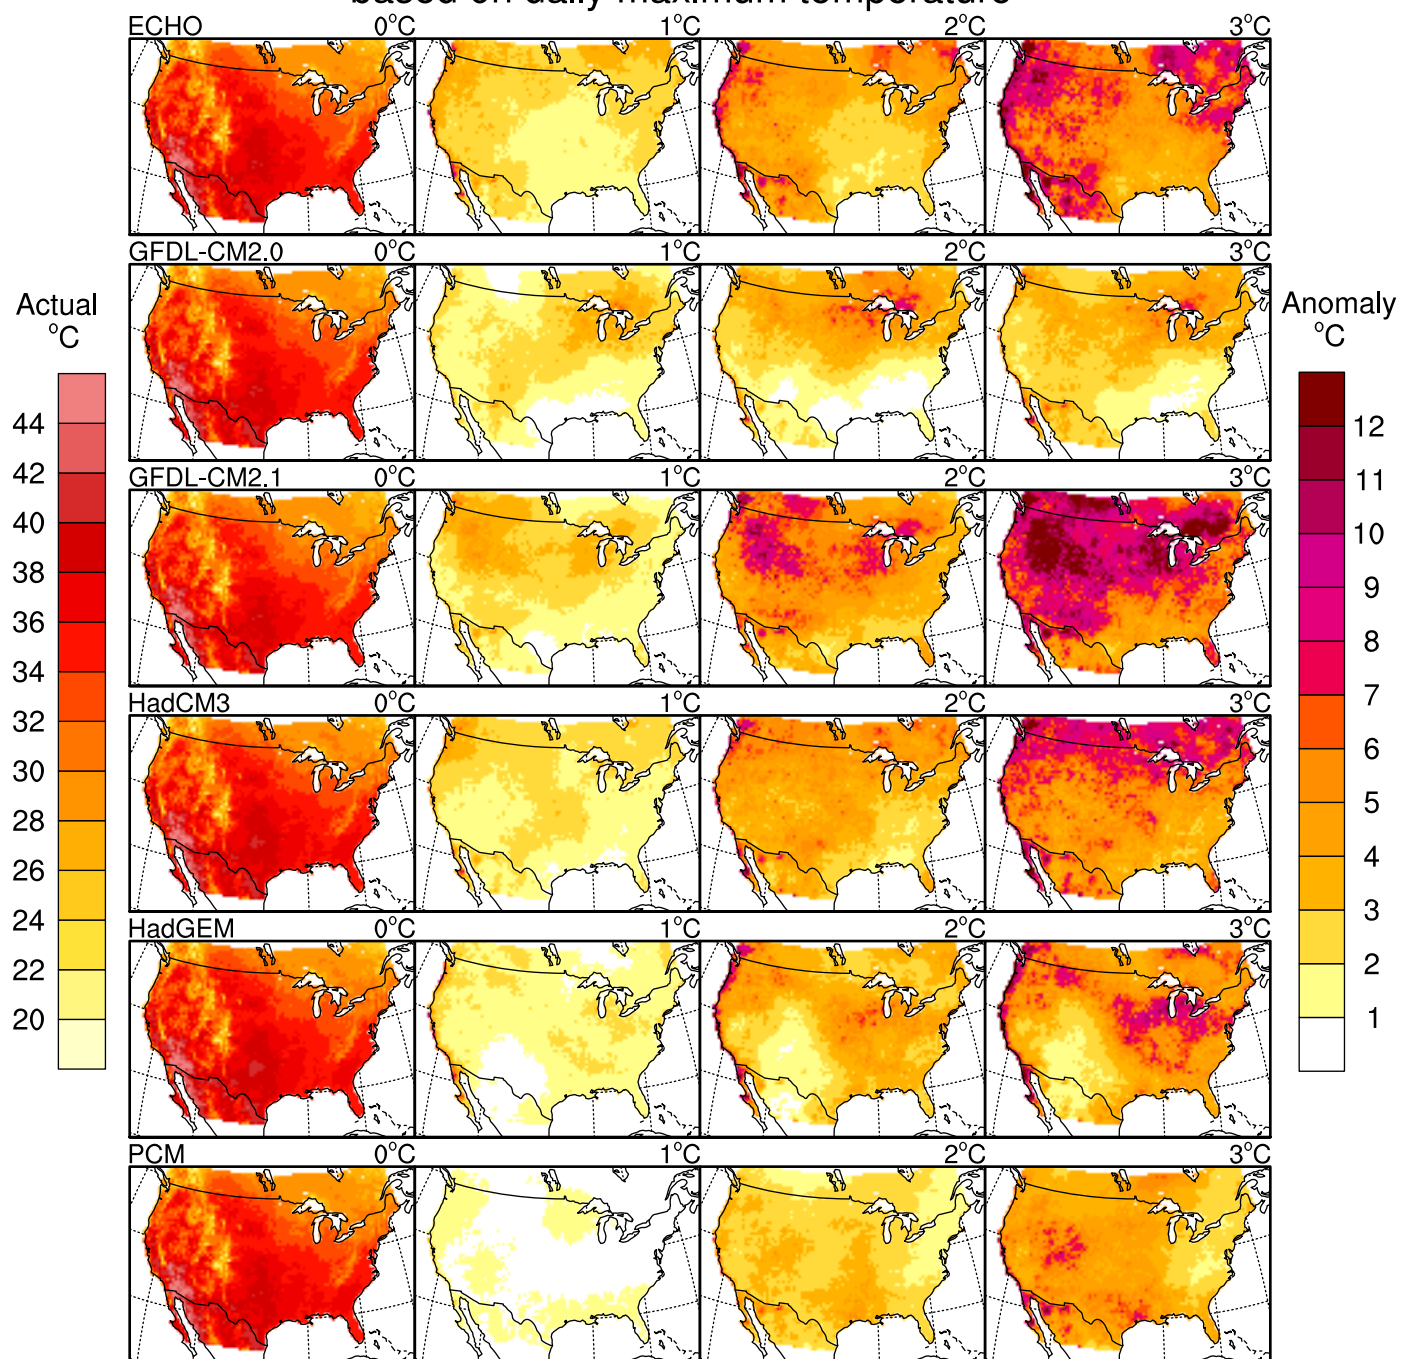
Temperature of the Hottest 3-Day Period based on daily maximum temperature



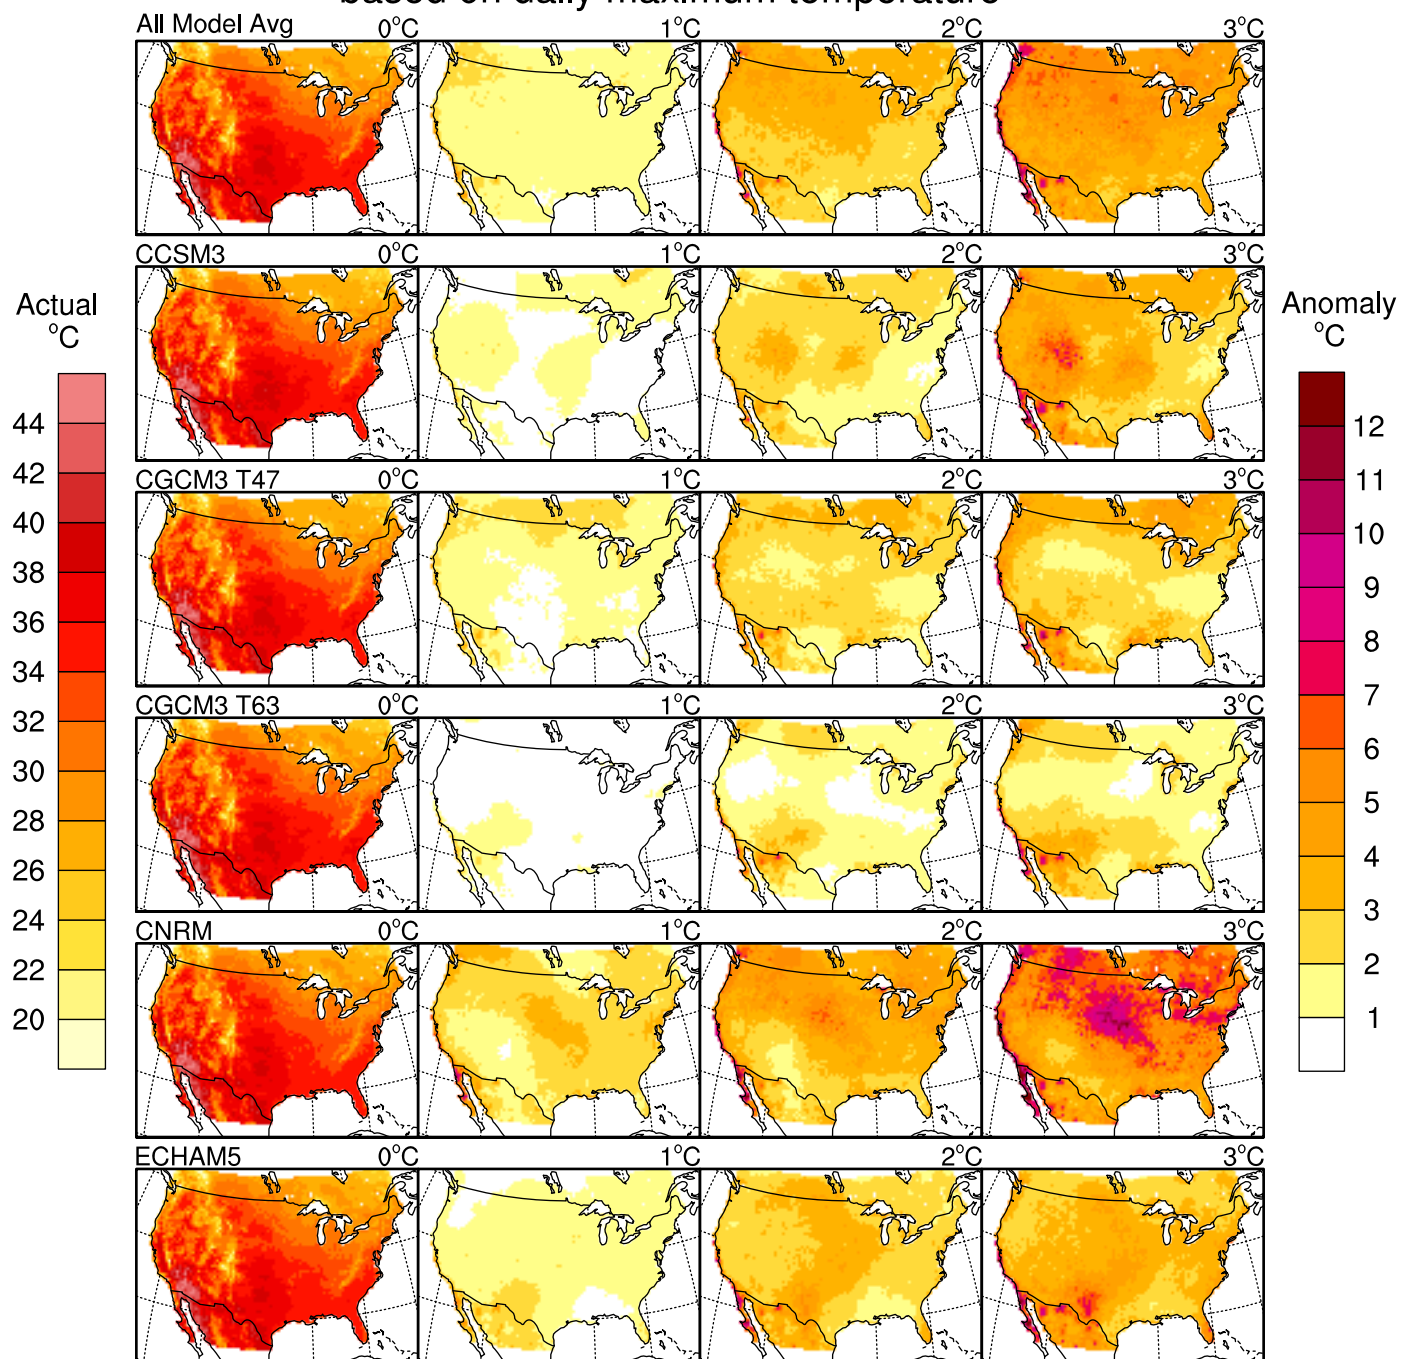
Temperature of the Hottest 5-Day Period based on daily maximum temperature



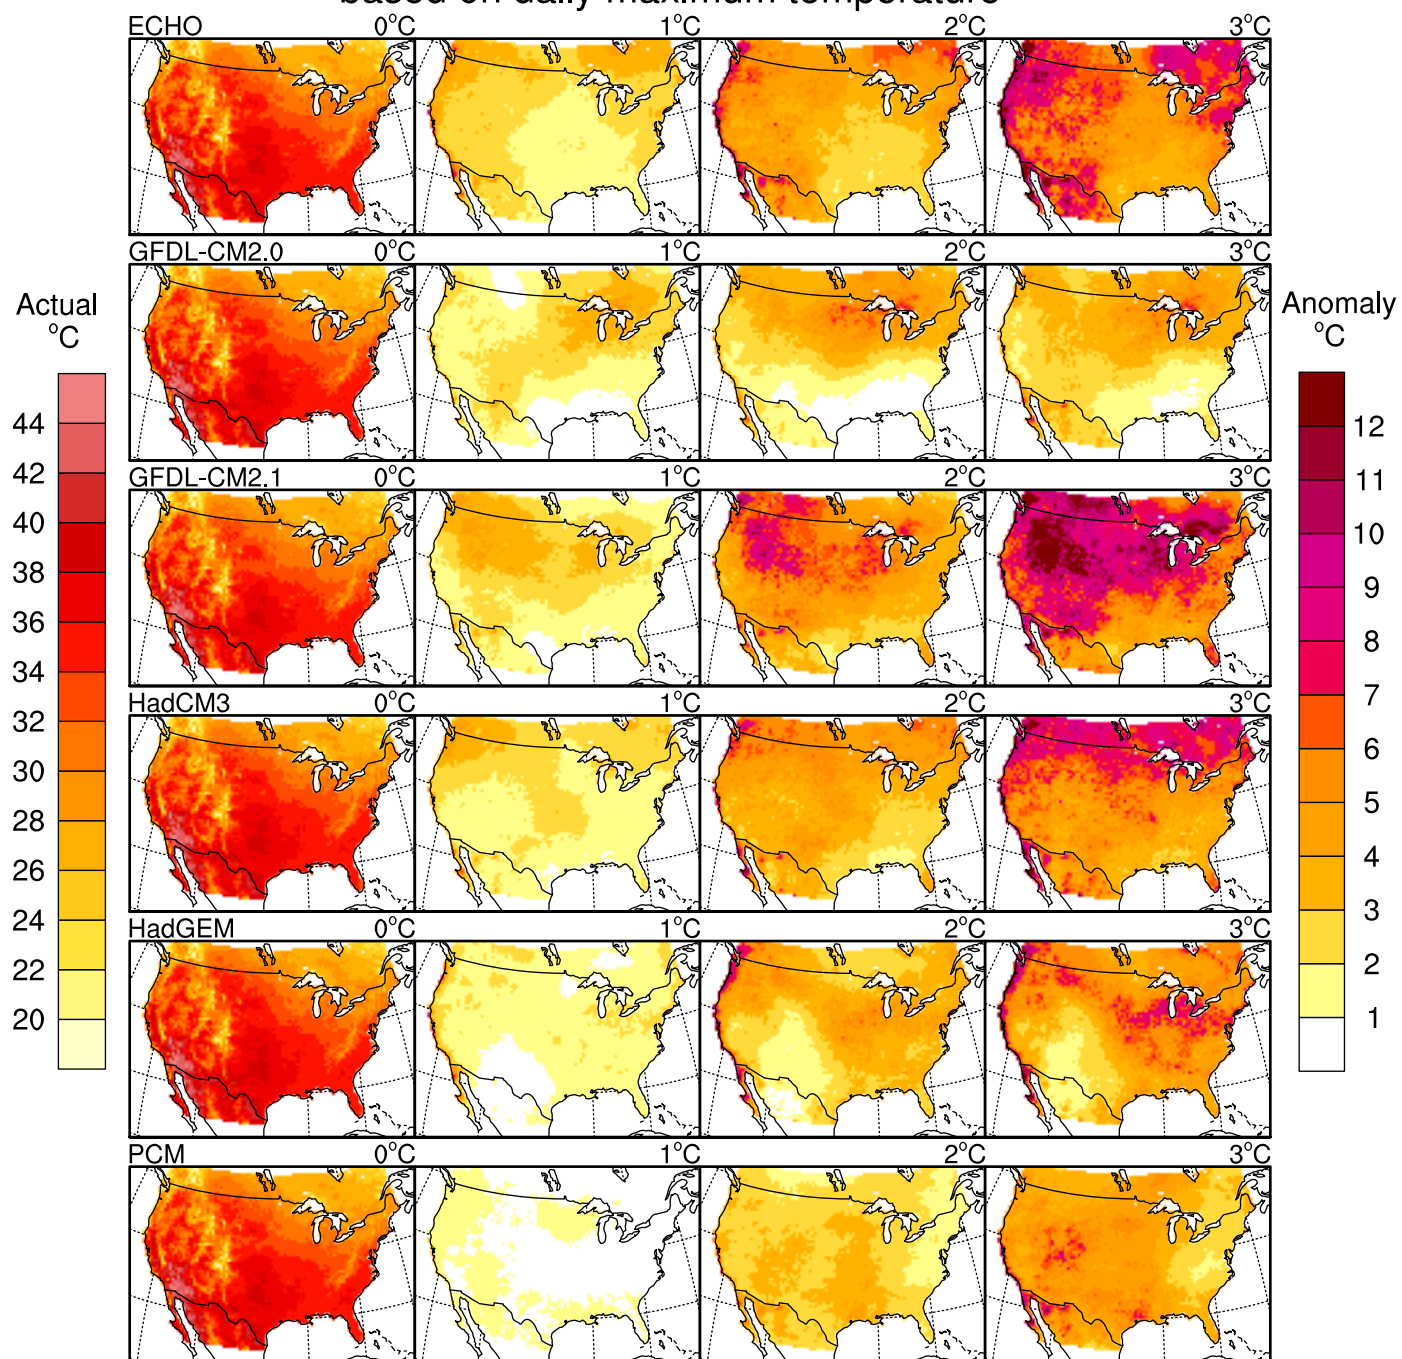
Temperature of the Hottest 5-Day Period based on daily maximum temperature



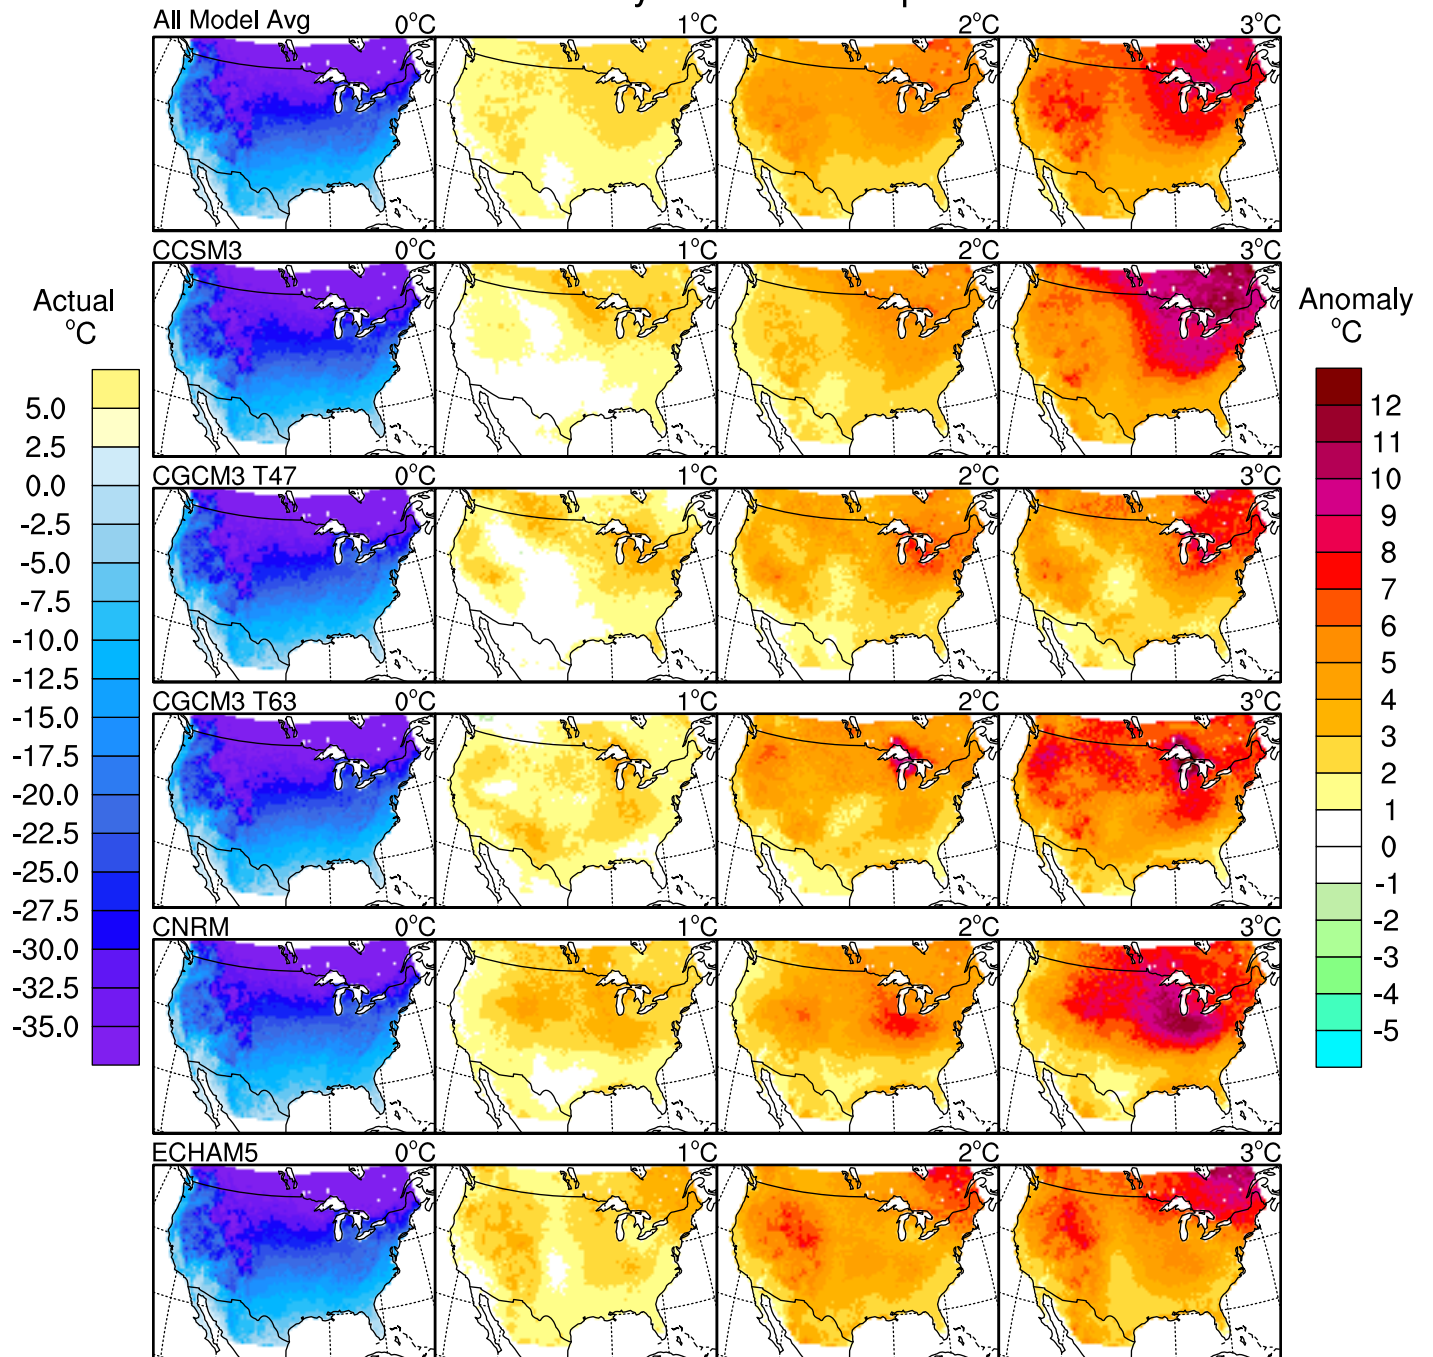
Temperature of the Hottest 10-Day Period based on daily maximum temperature



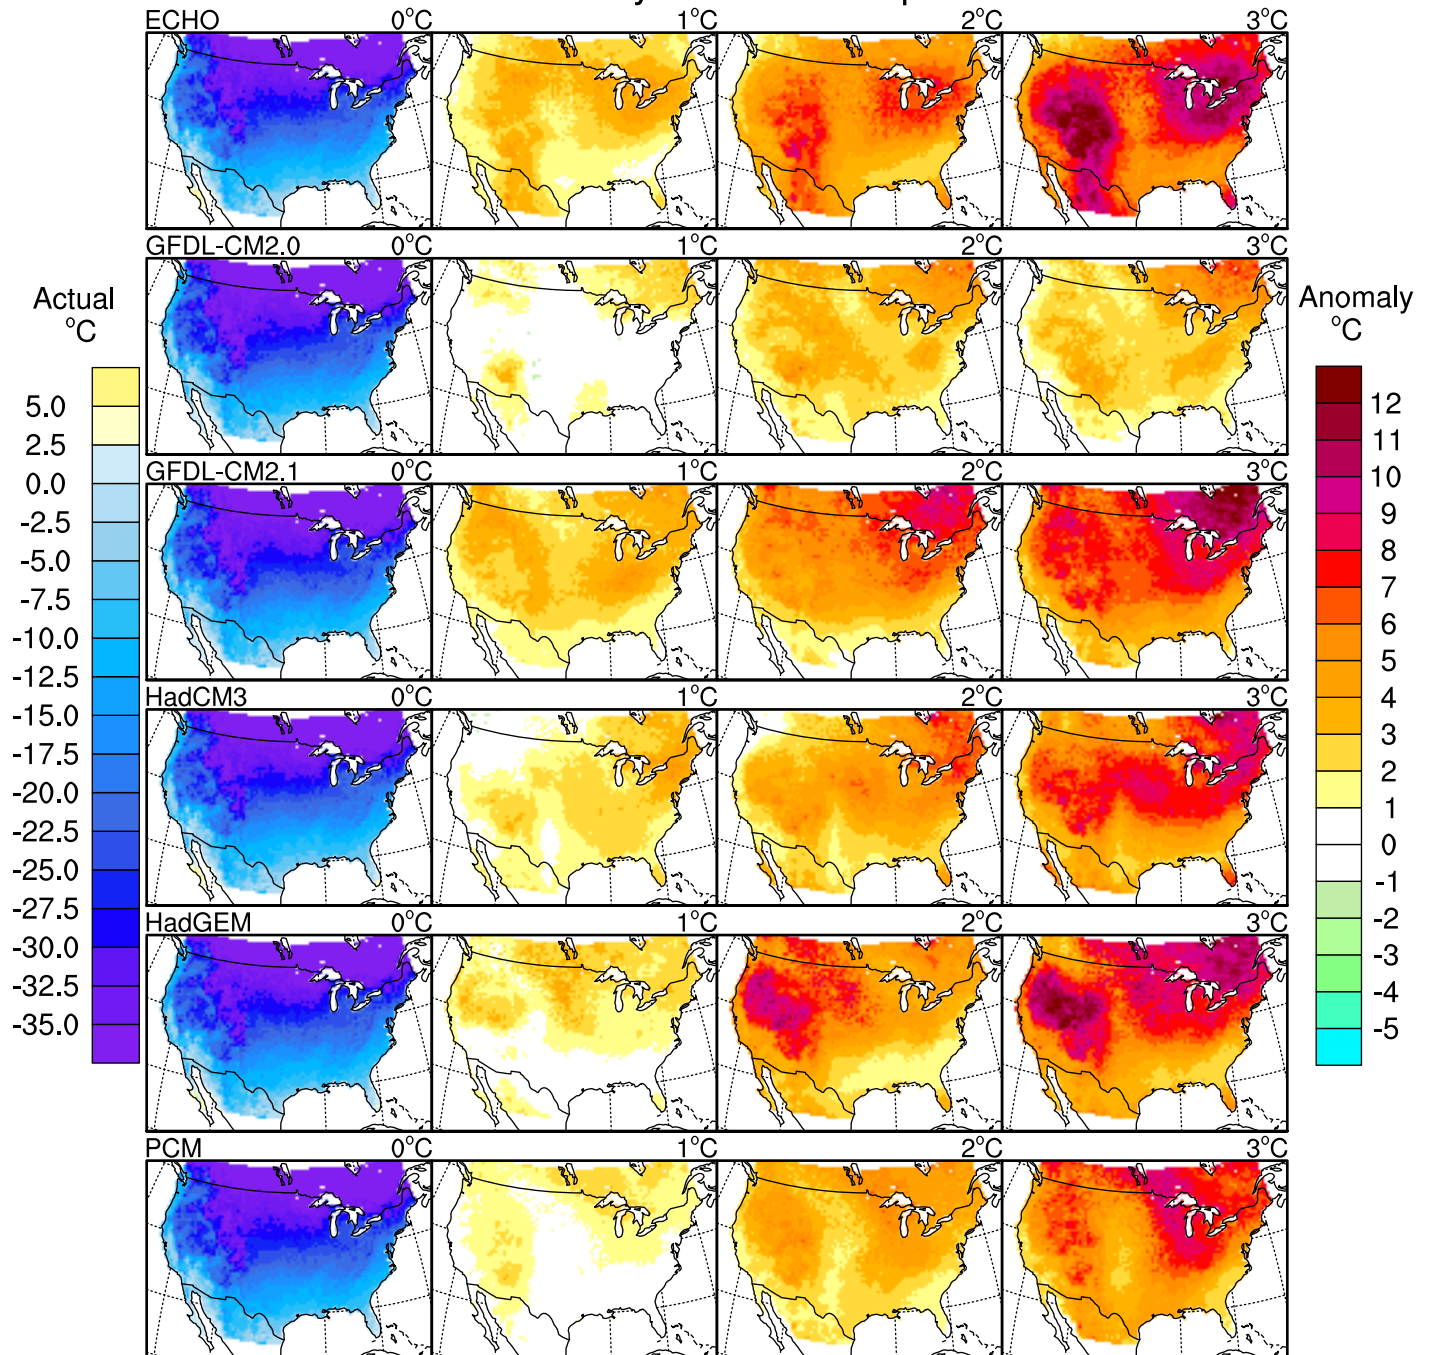
Temperature of the Hottest 10-Day Period based on daily maximum temperature



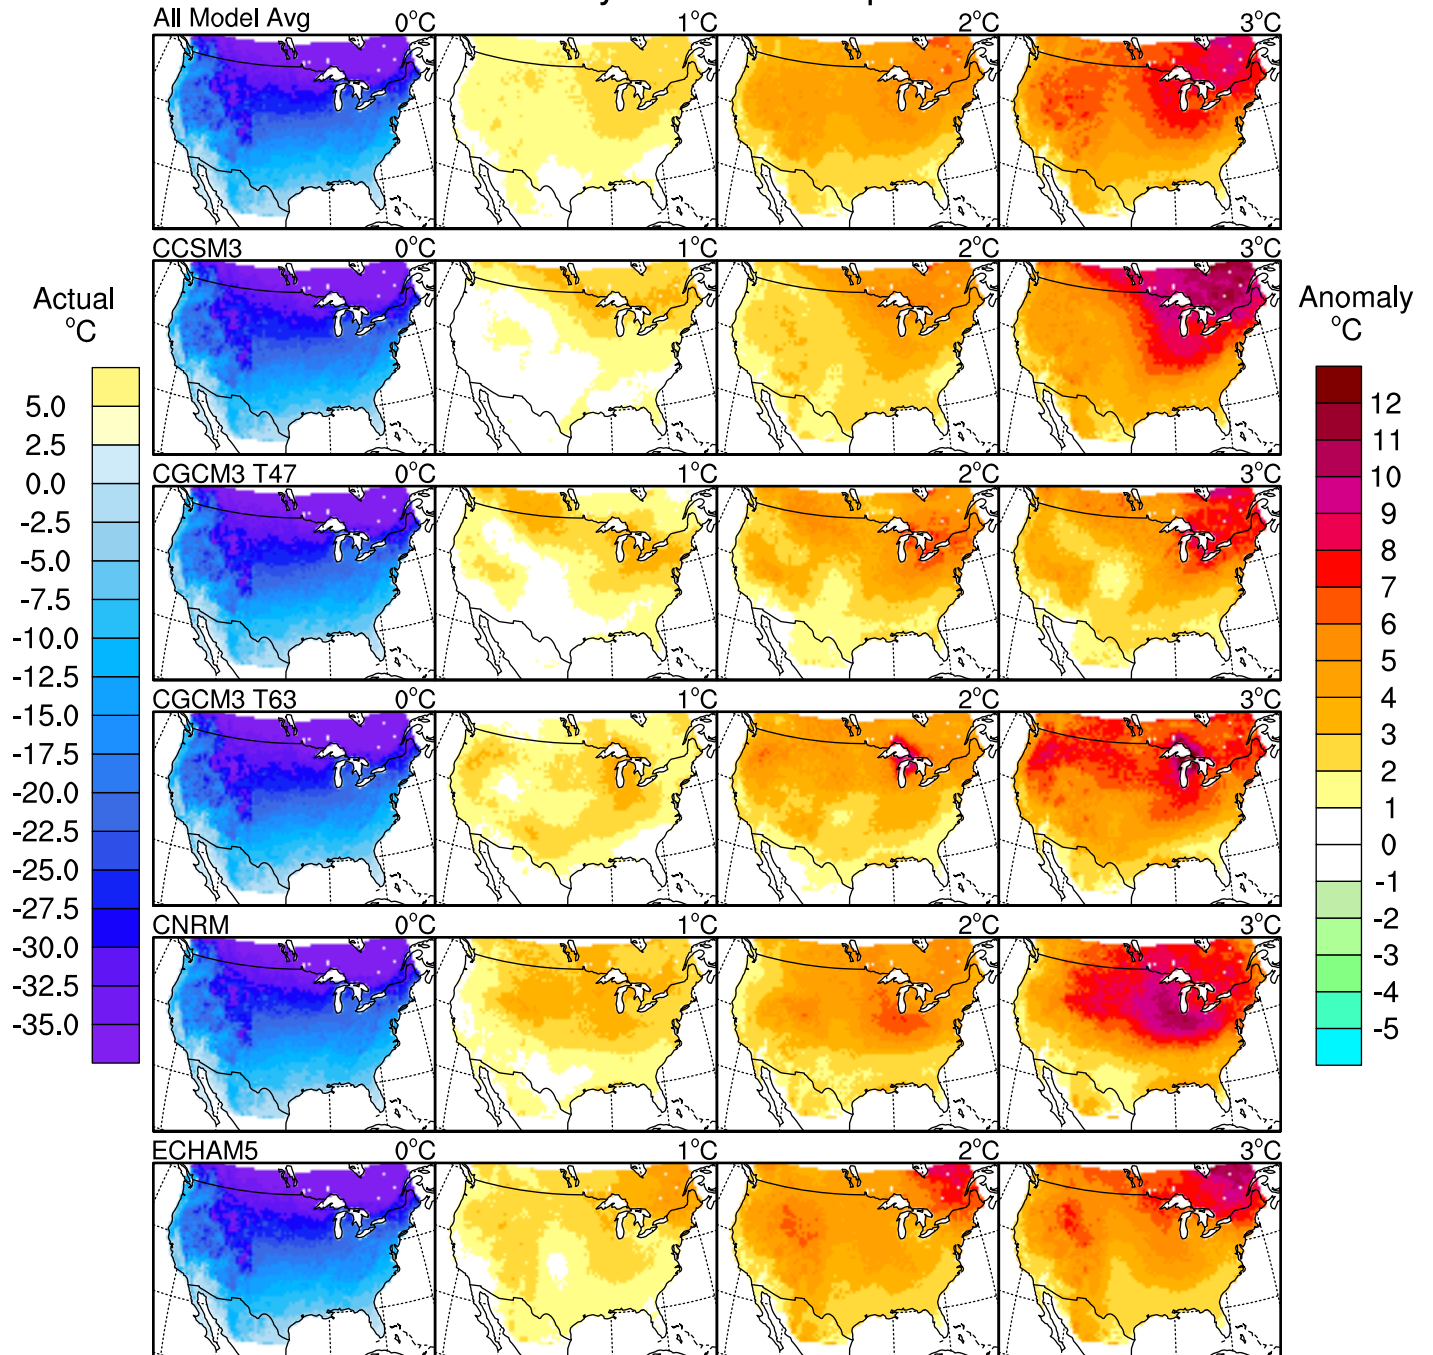
Temperature of the Coldest Day based on daily minimum temperature



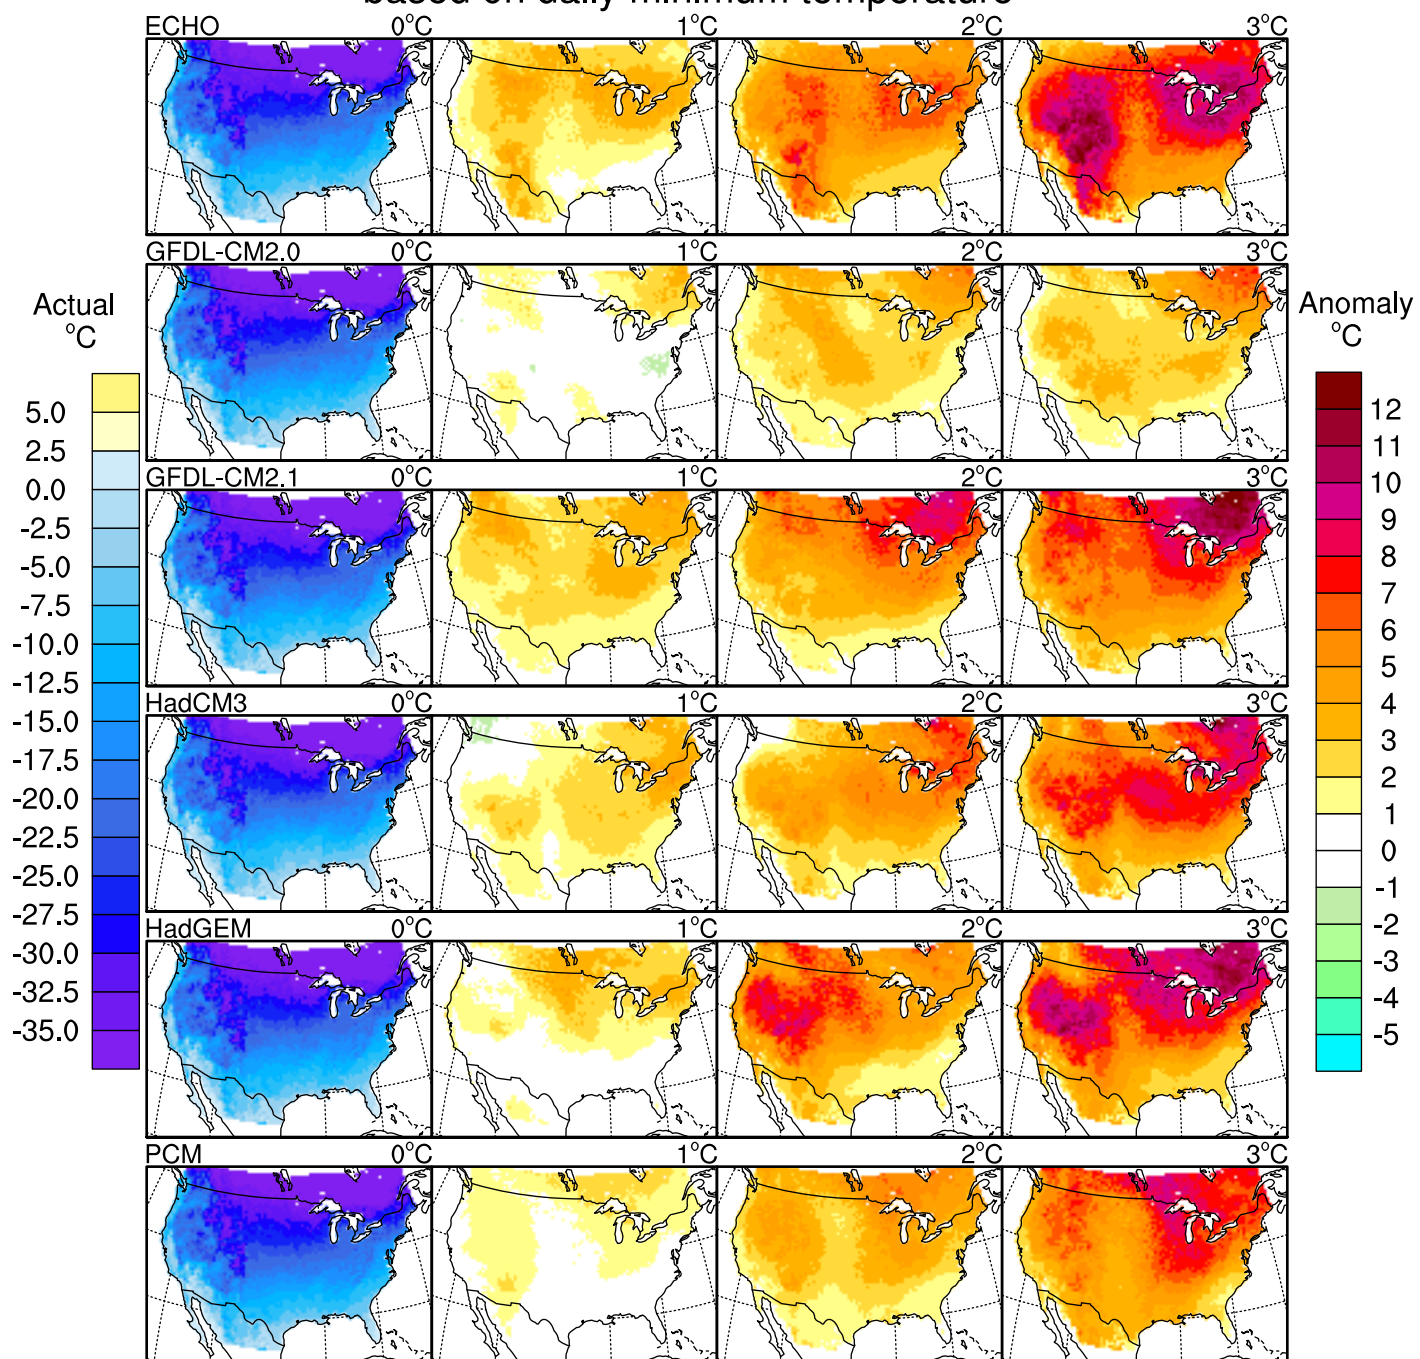
Temperature of the Coldest Day based on daily minimum temperature



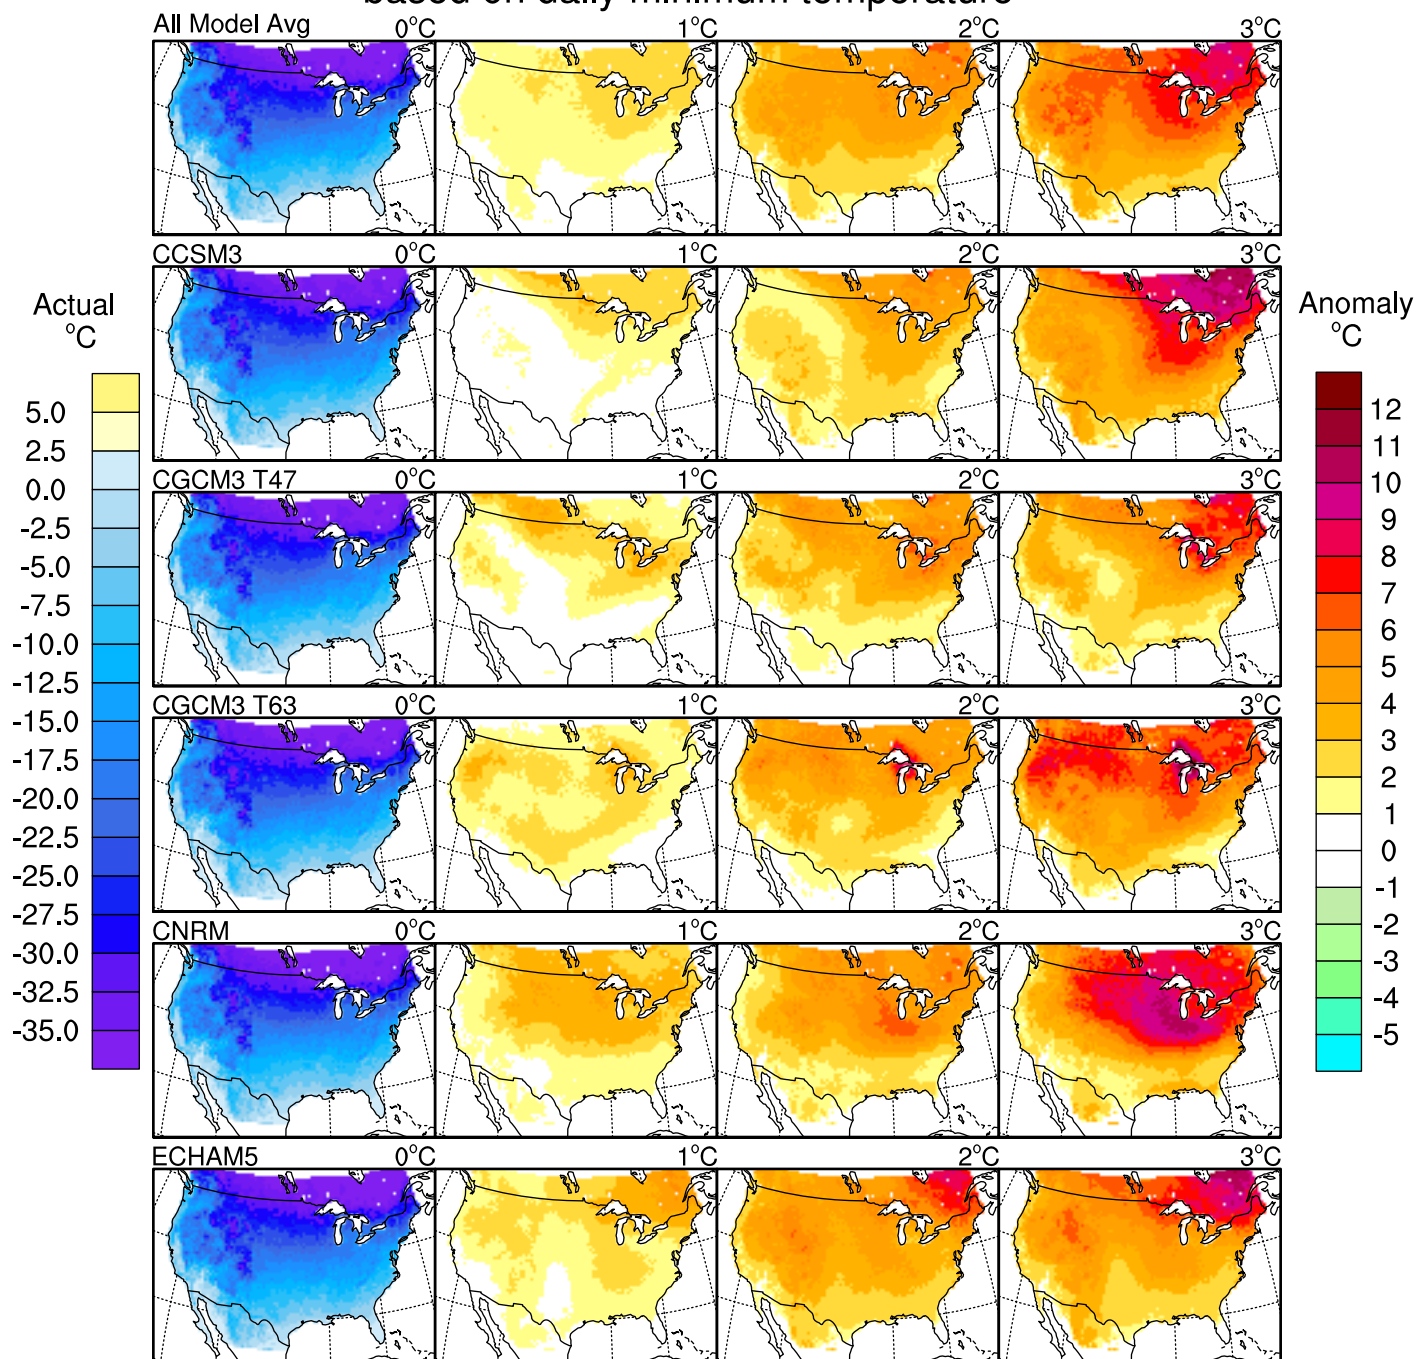
Temperature of the Coldest 3-Day Period based on daily minimum temperature



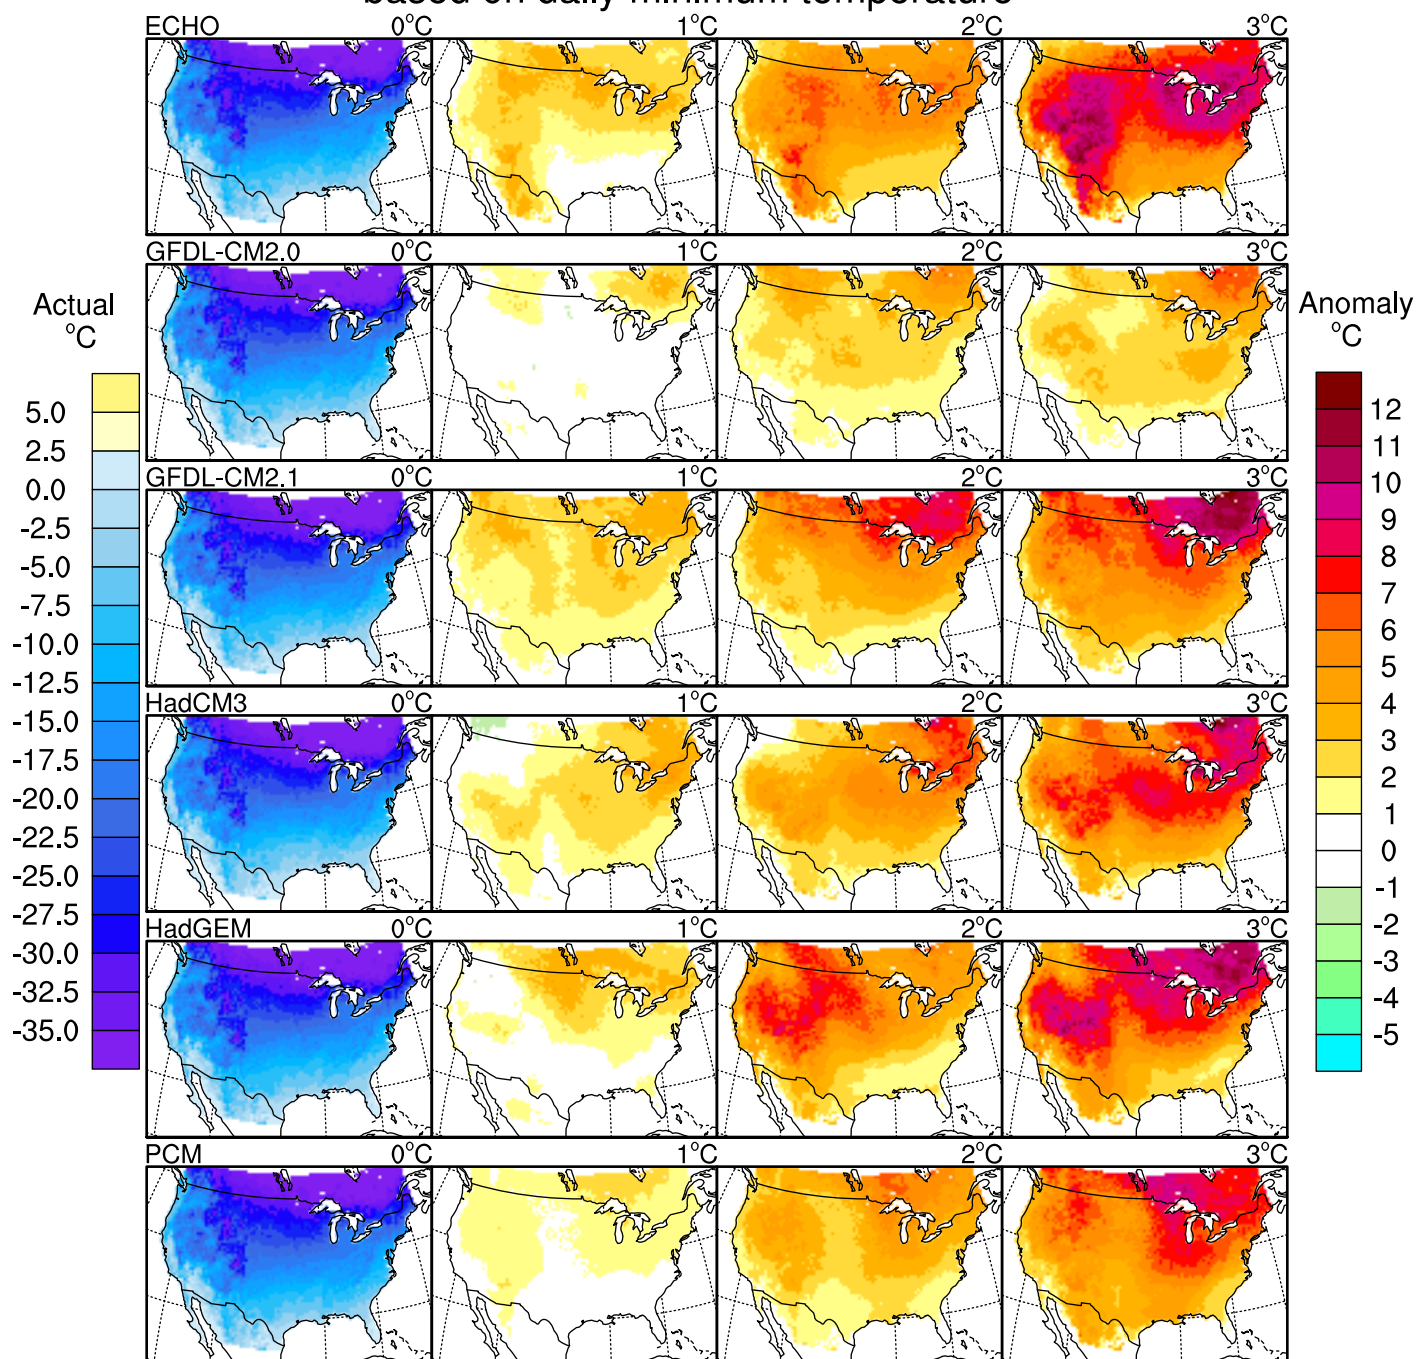
Temperature of the Coldest 3-Day Period based on daily minimum temperature



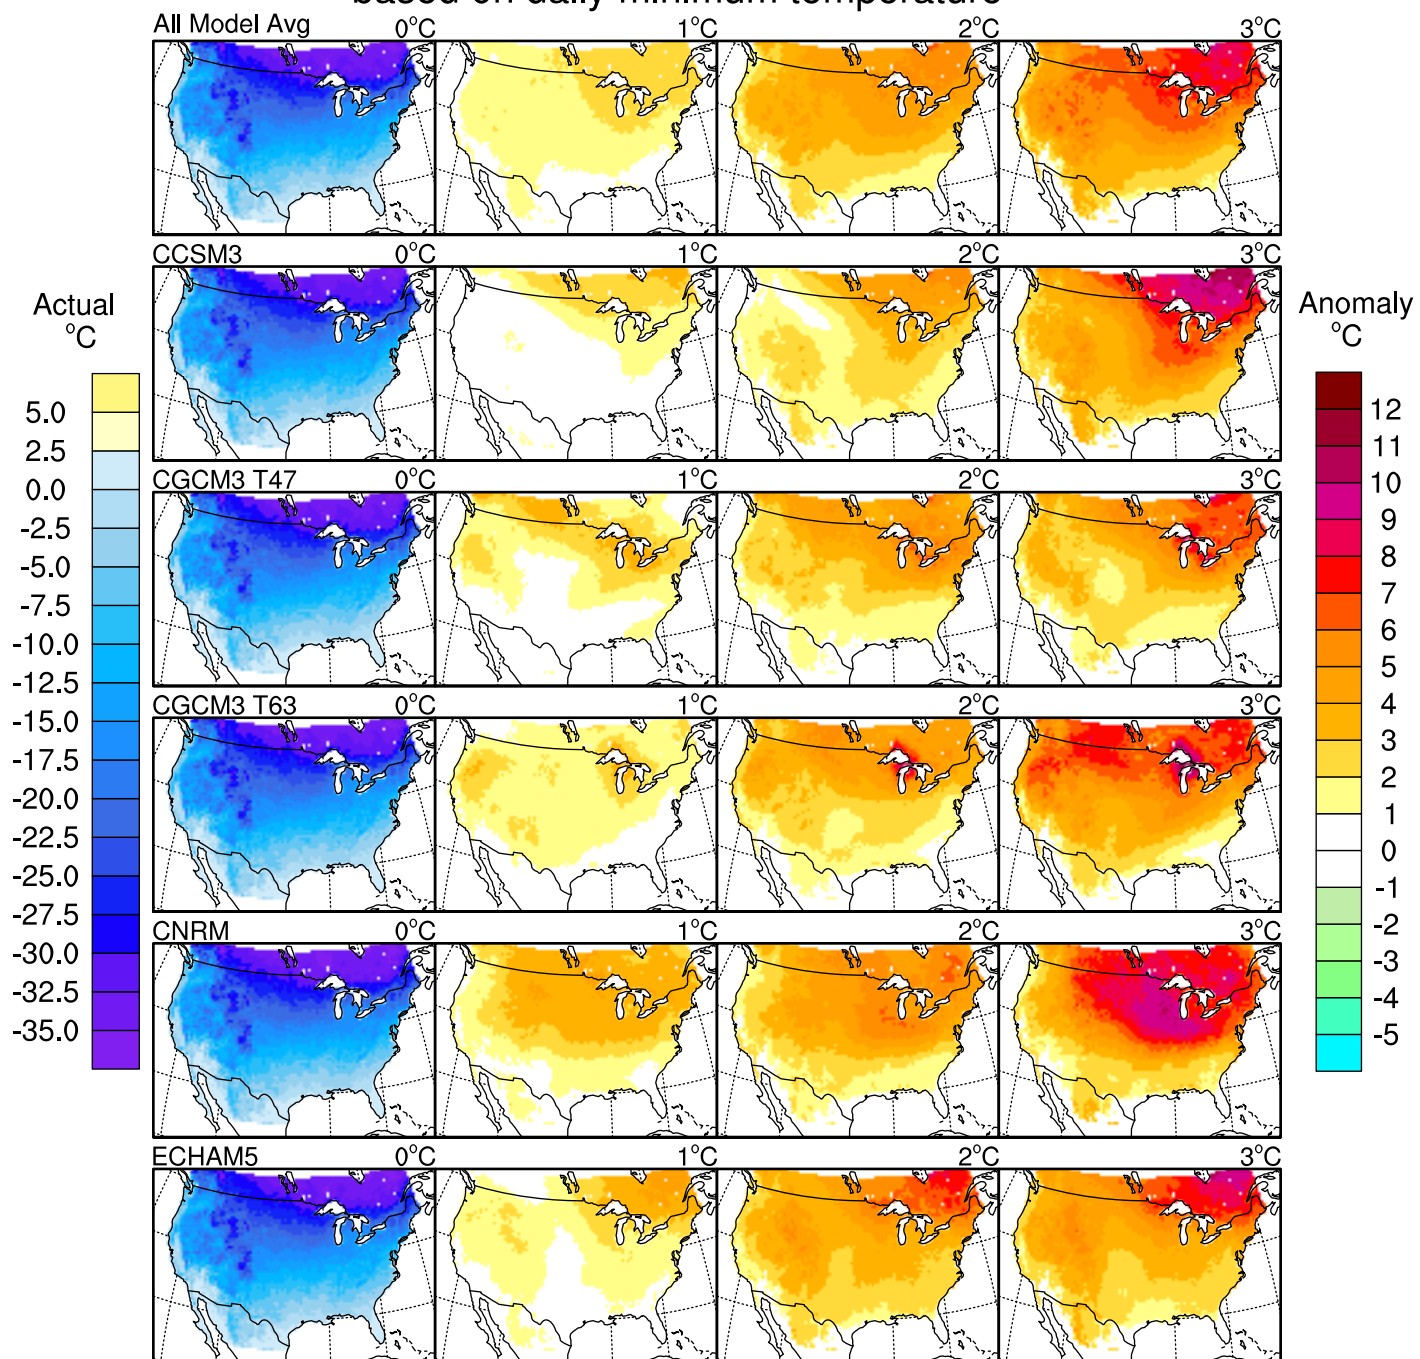
Temperature of the Coldest 5-Day Period based on daily minimum temperature



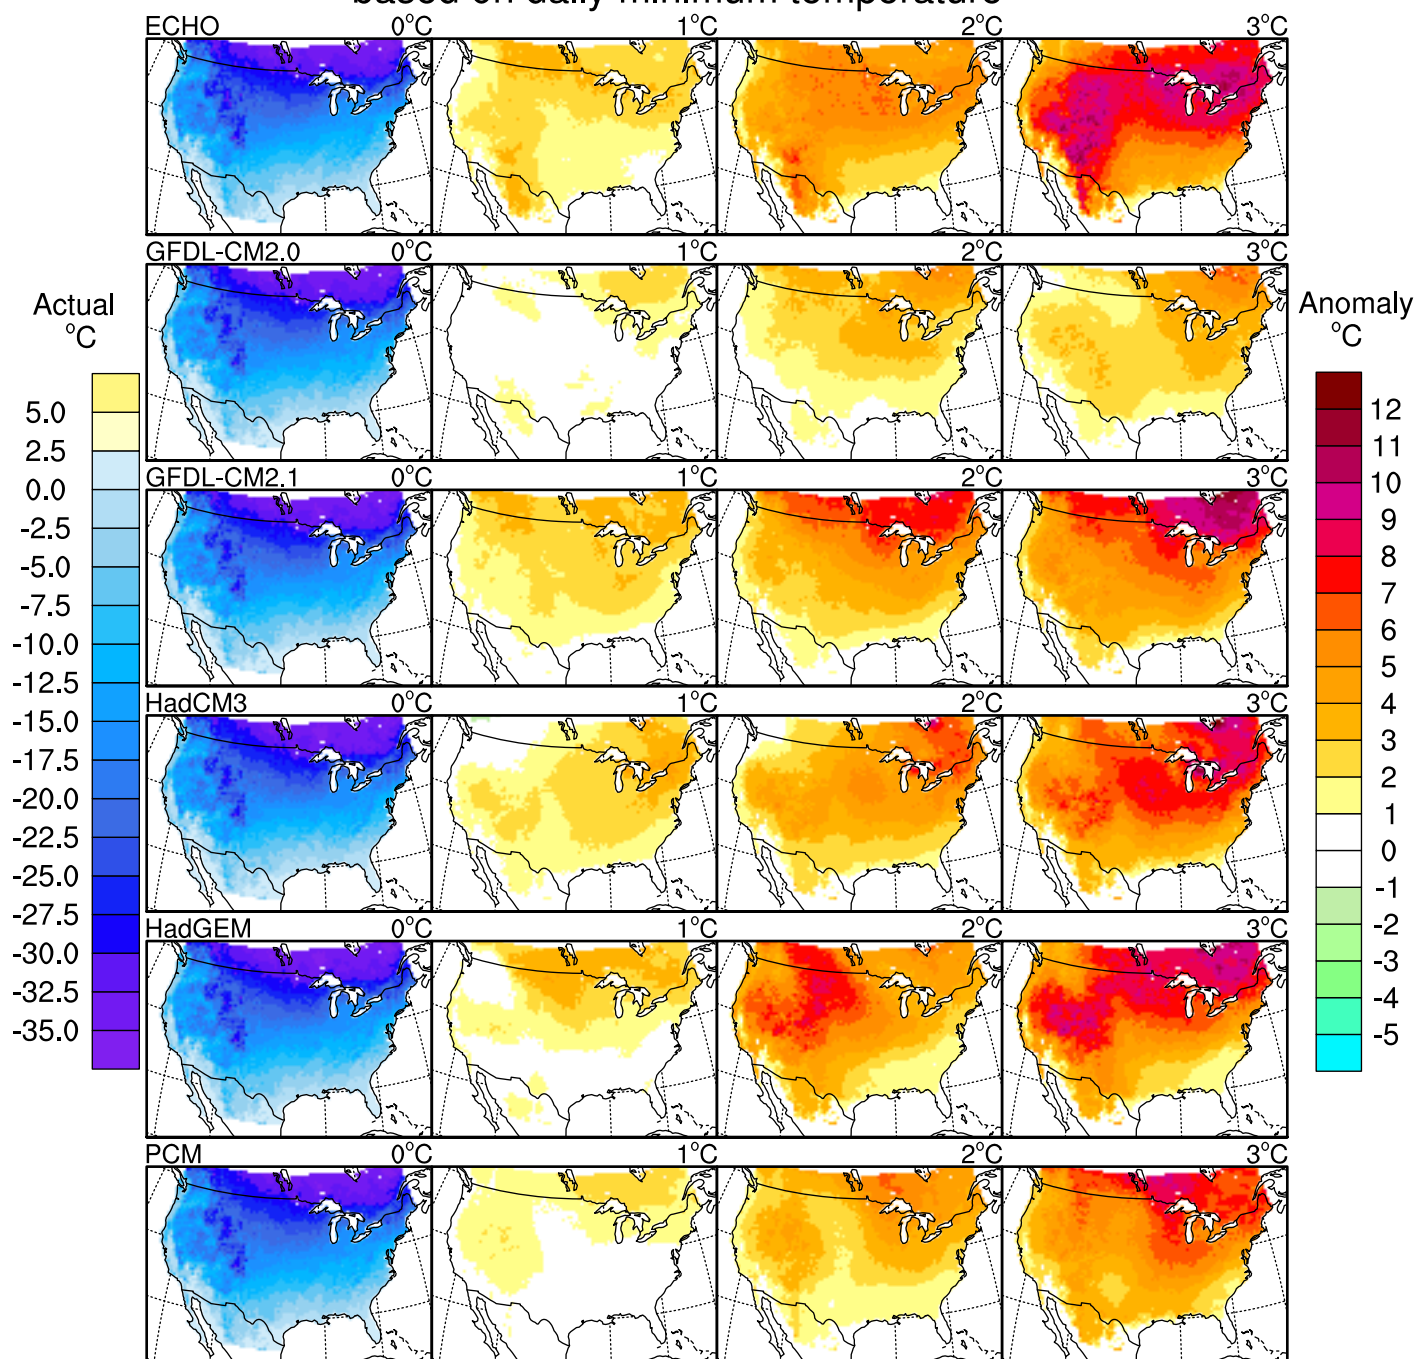
Temperature of the Coldest 5-Day Period based on daily minimum temperature



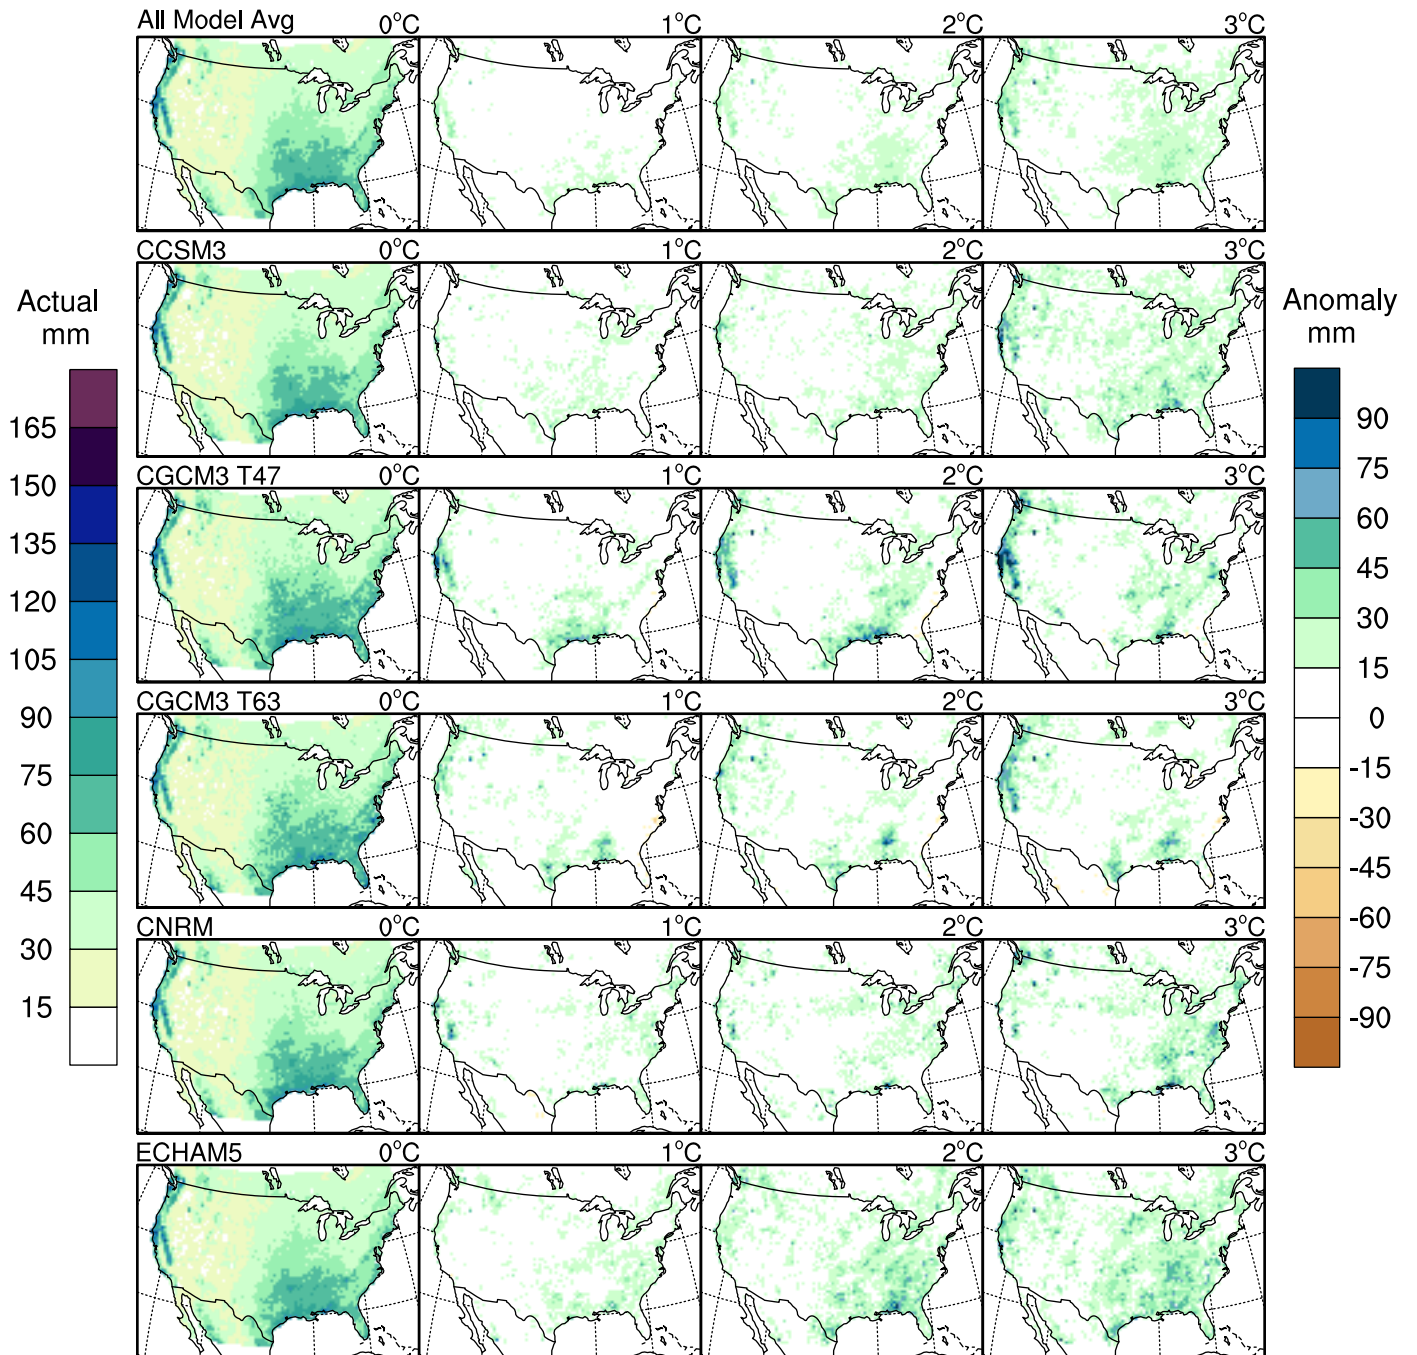
Temperature of the Coldest 10-Day Period based on daily minimum temperature



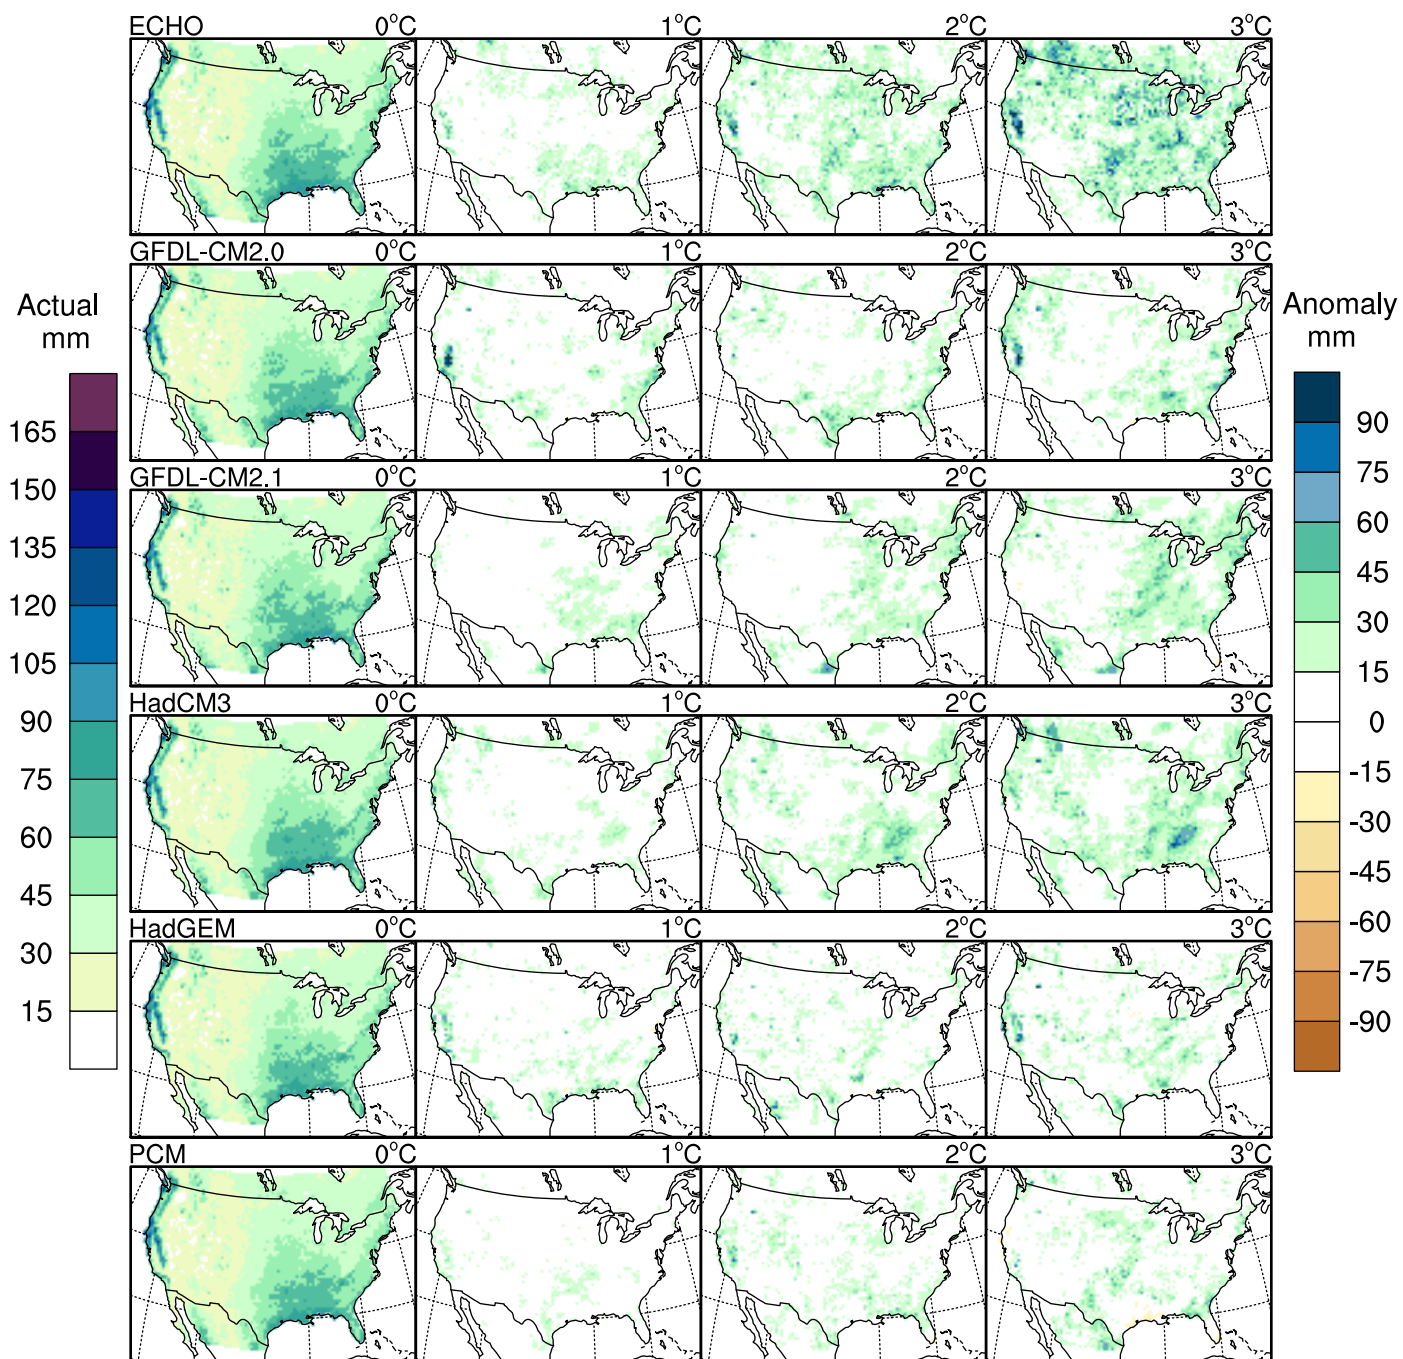
Temperature of the Coldest 10-Day Period based on daily minimum temperature



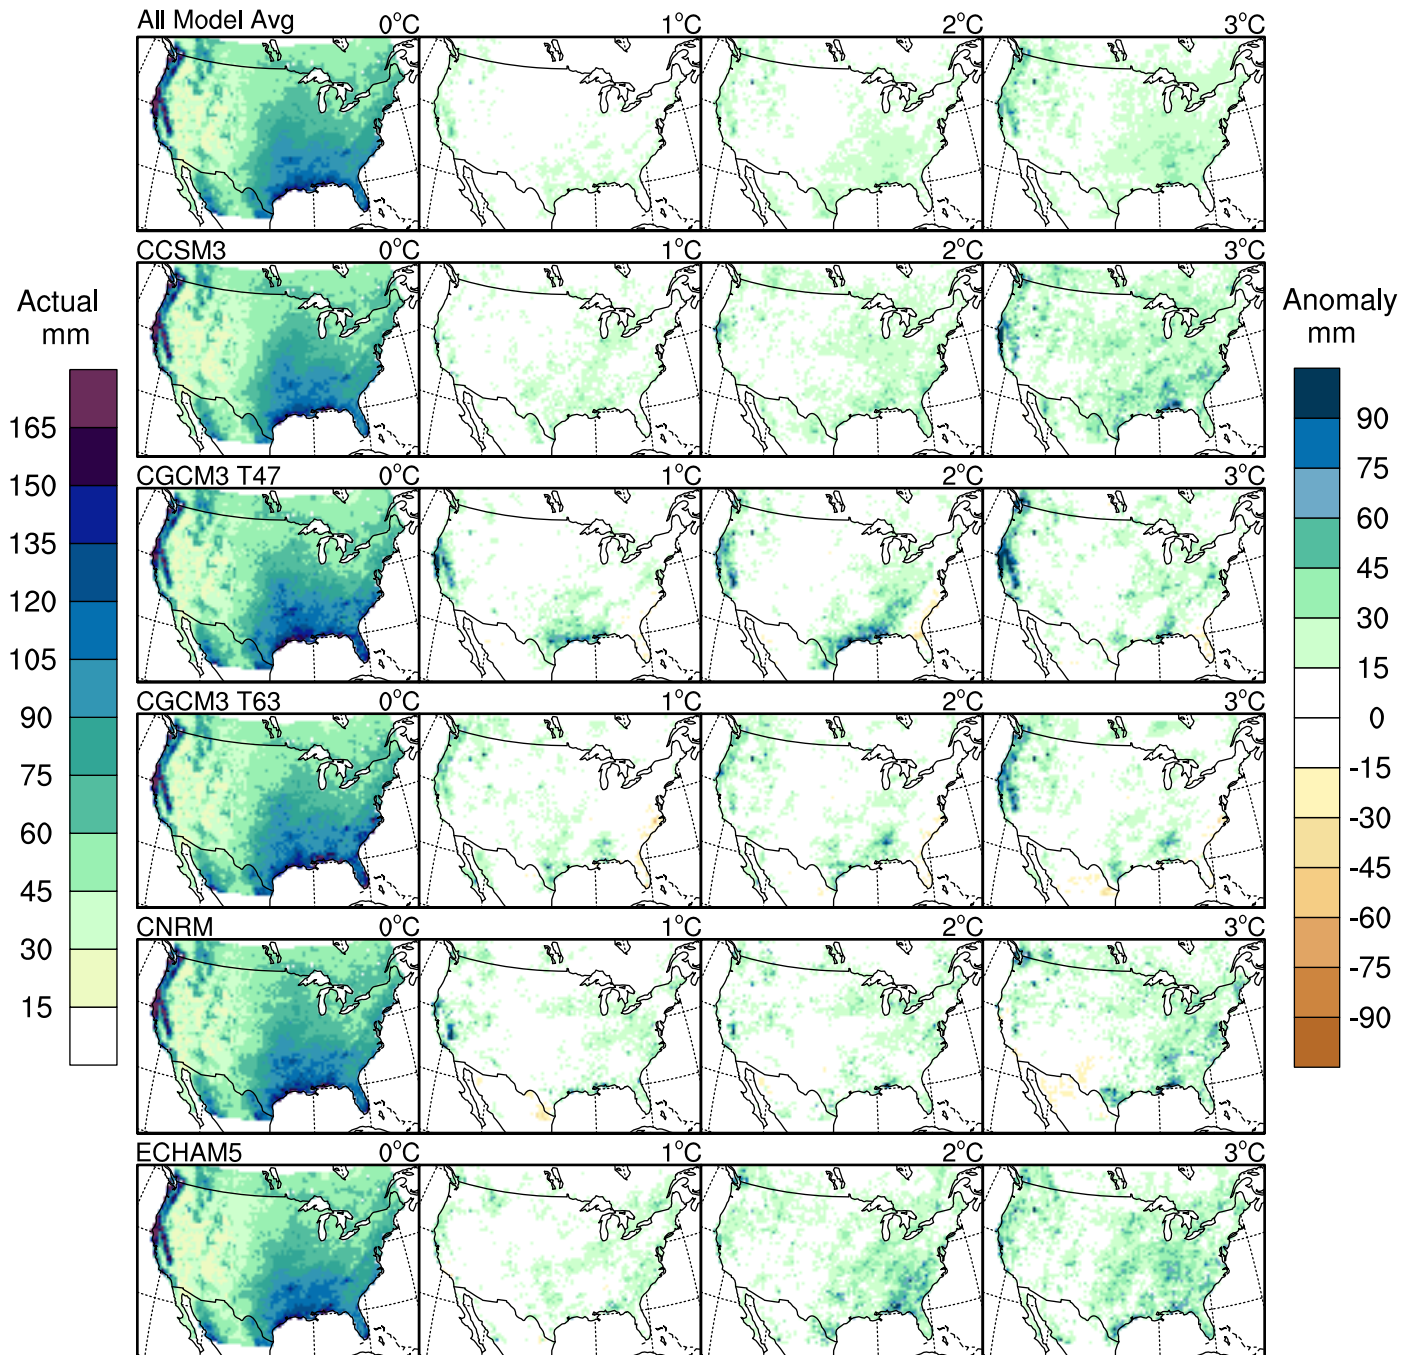
Cumulative Precipitation of the Wettest Day



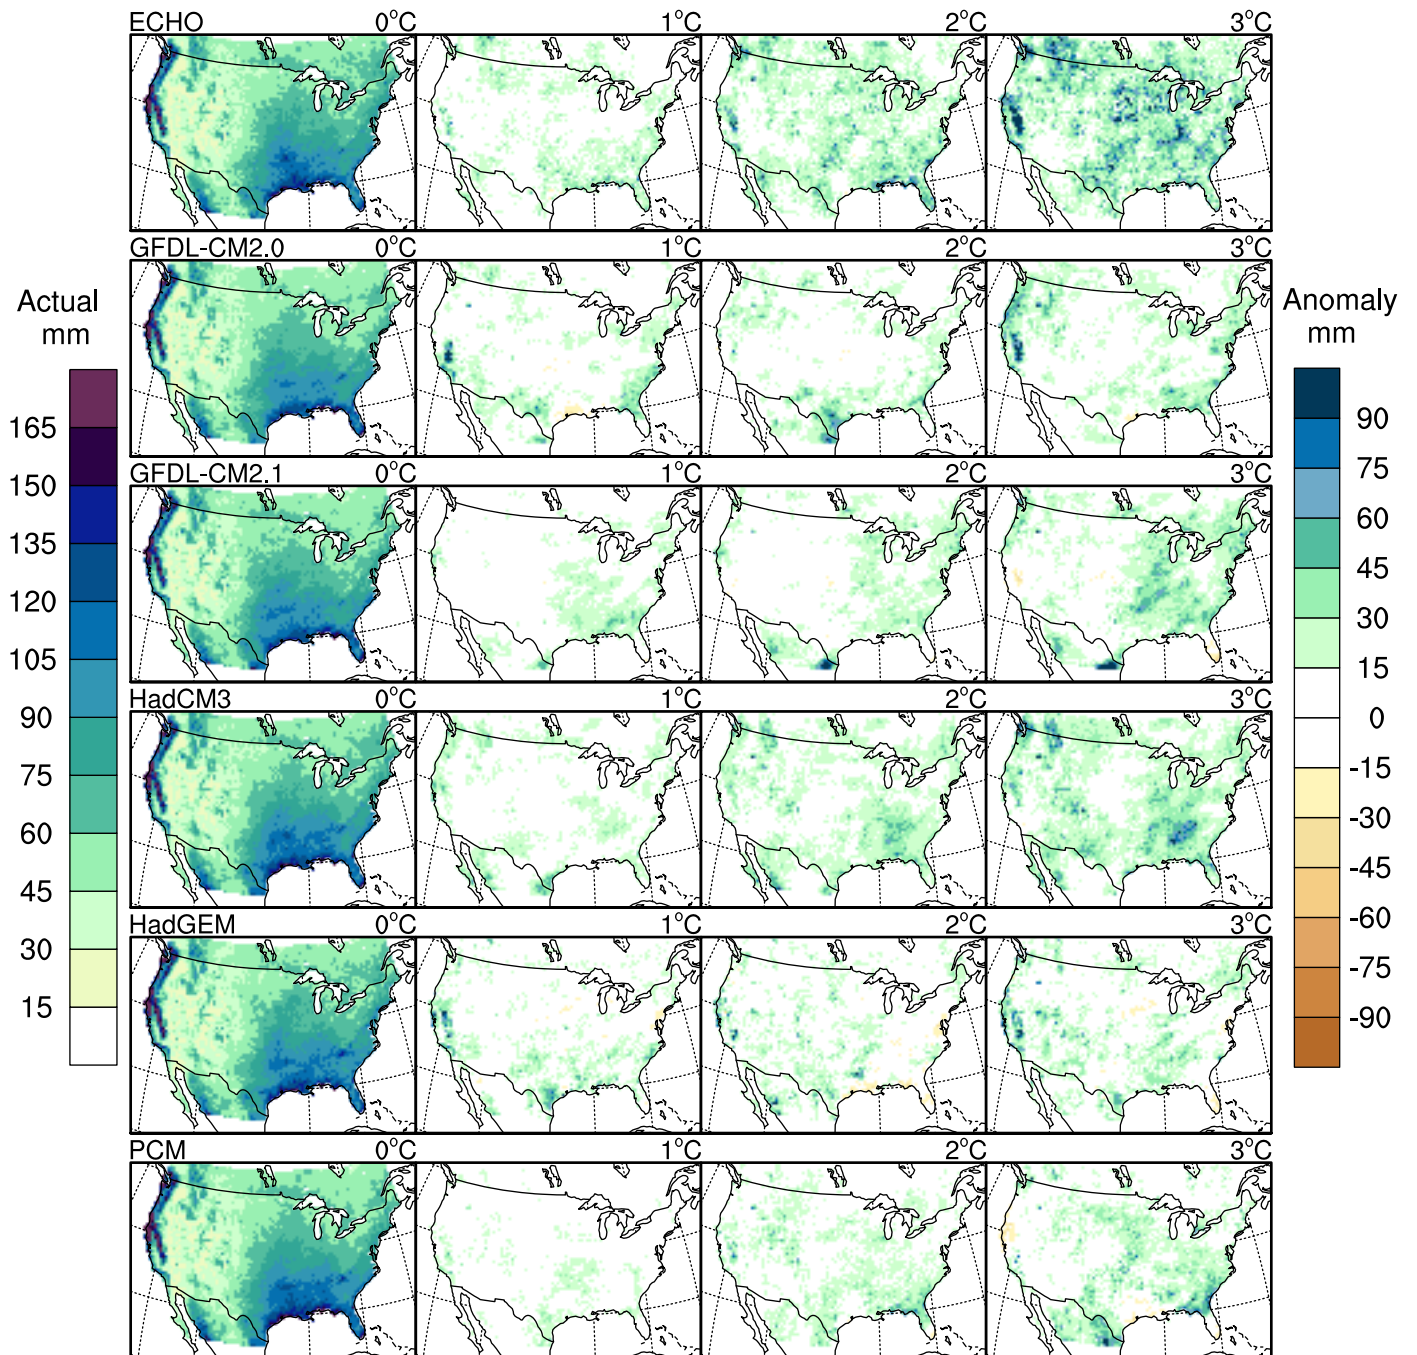
Cumulative Precipitation of the Wettest Day



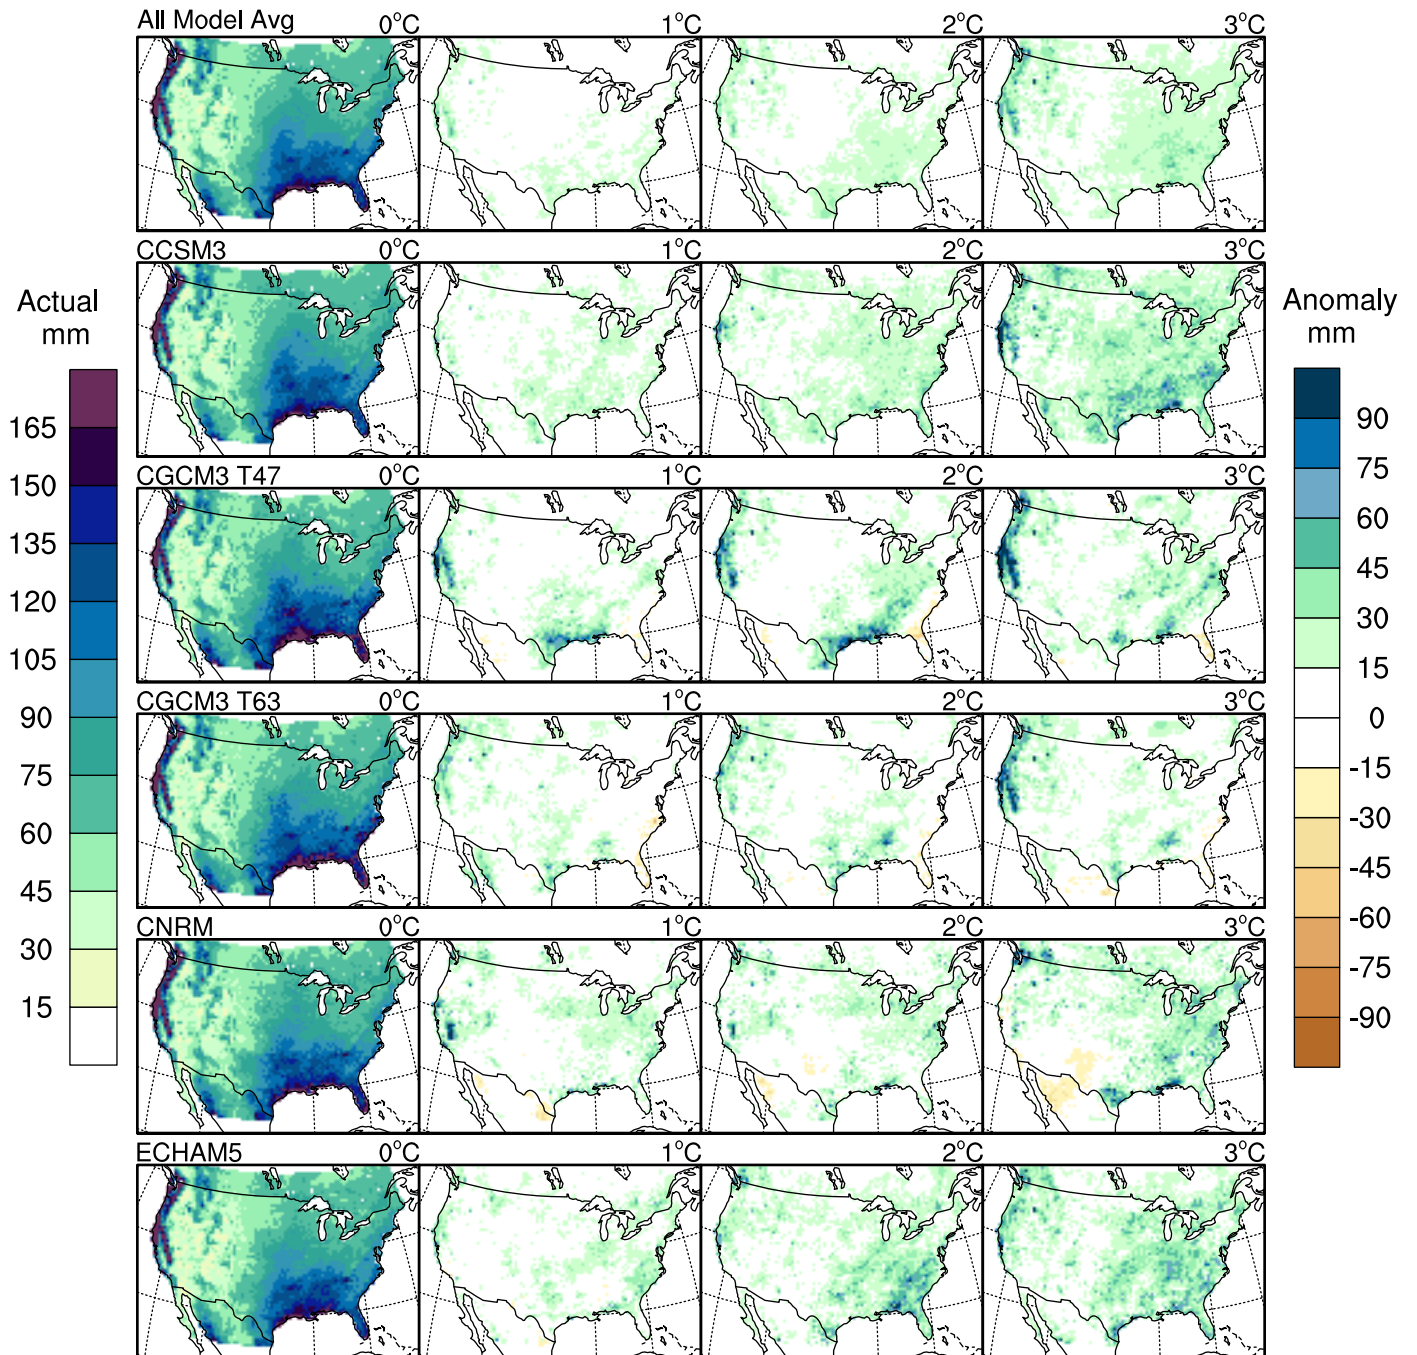
Cumulative Precipitation of the Wettest 3-Day Period



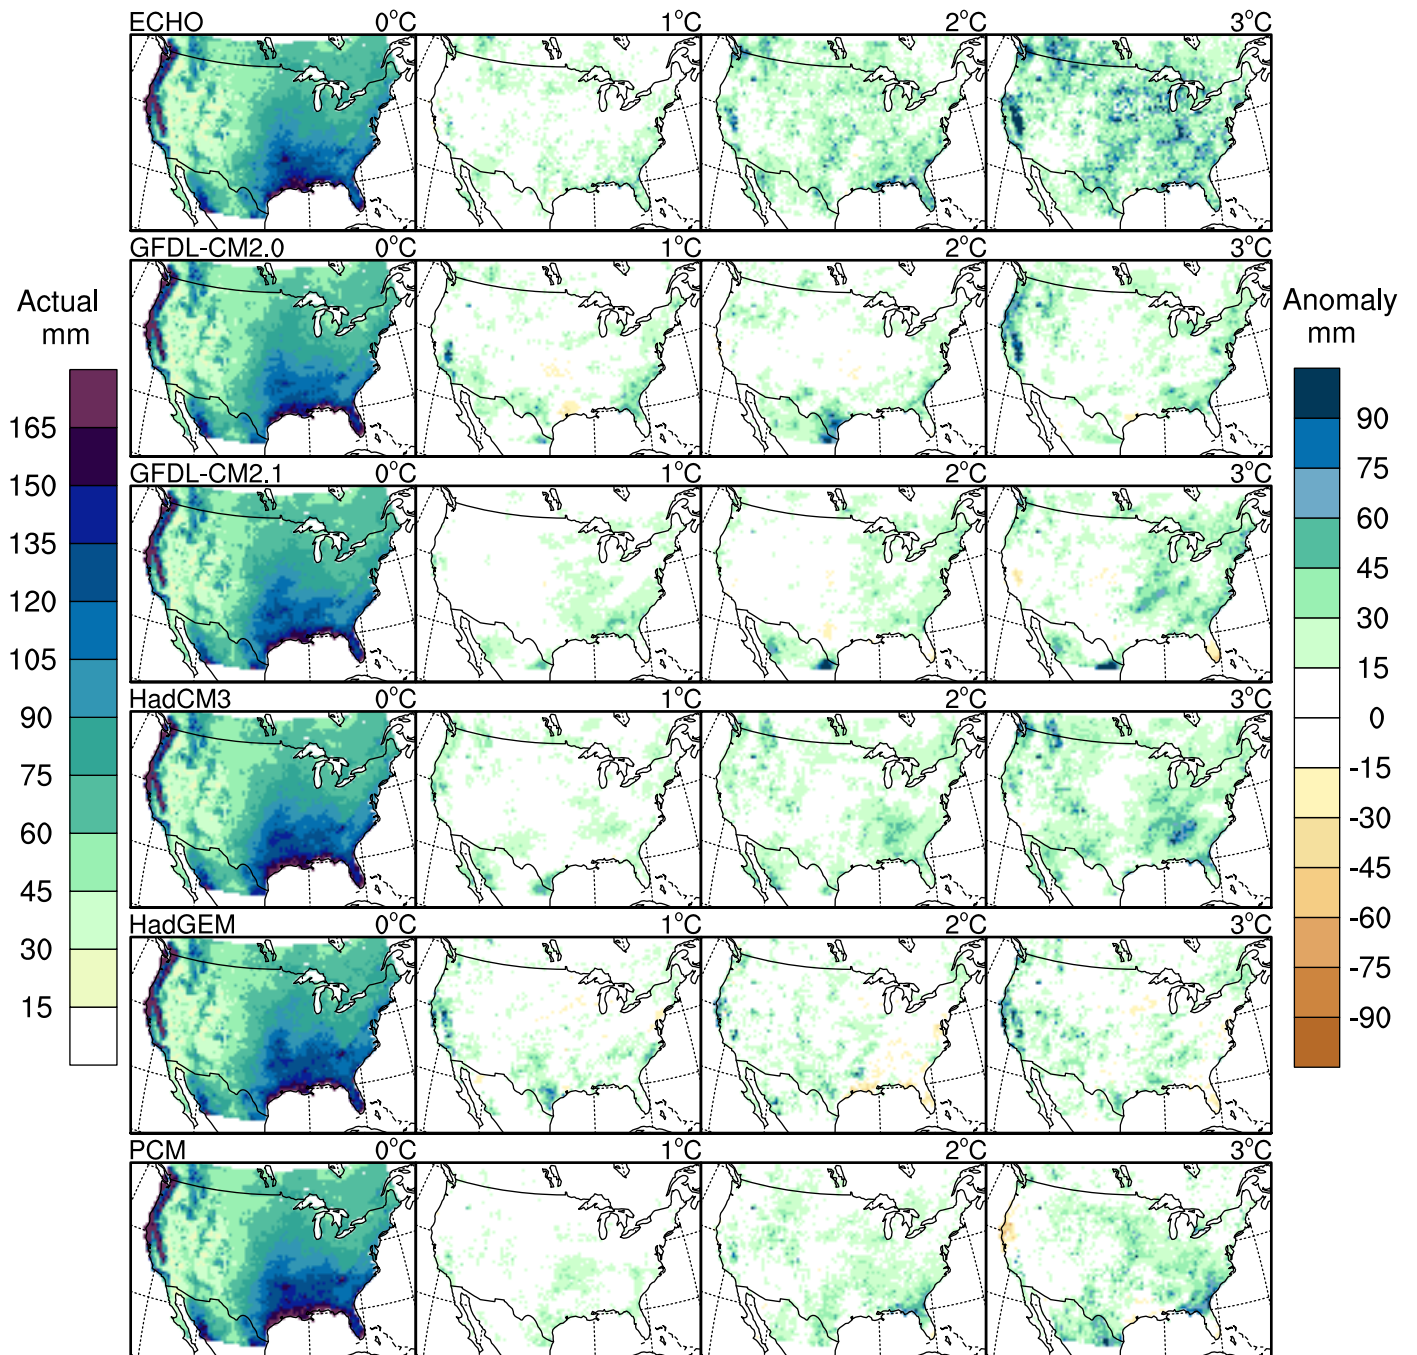
Cumulative Precipitation of the Wettest 3-Day Period



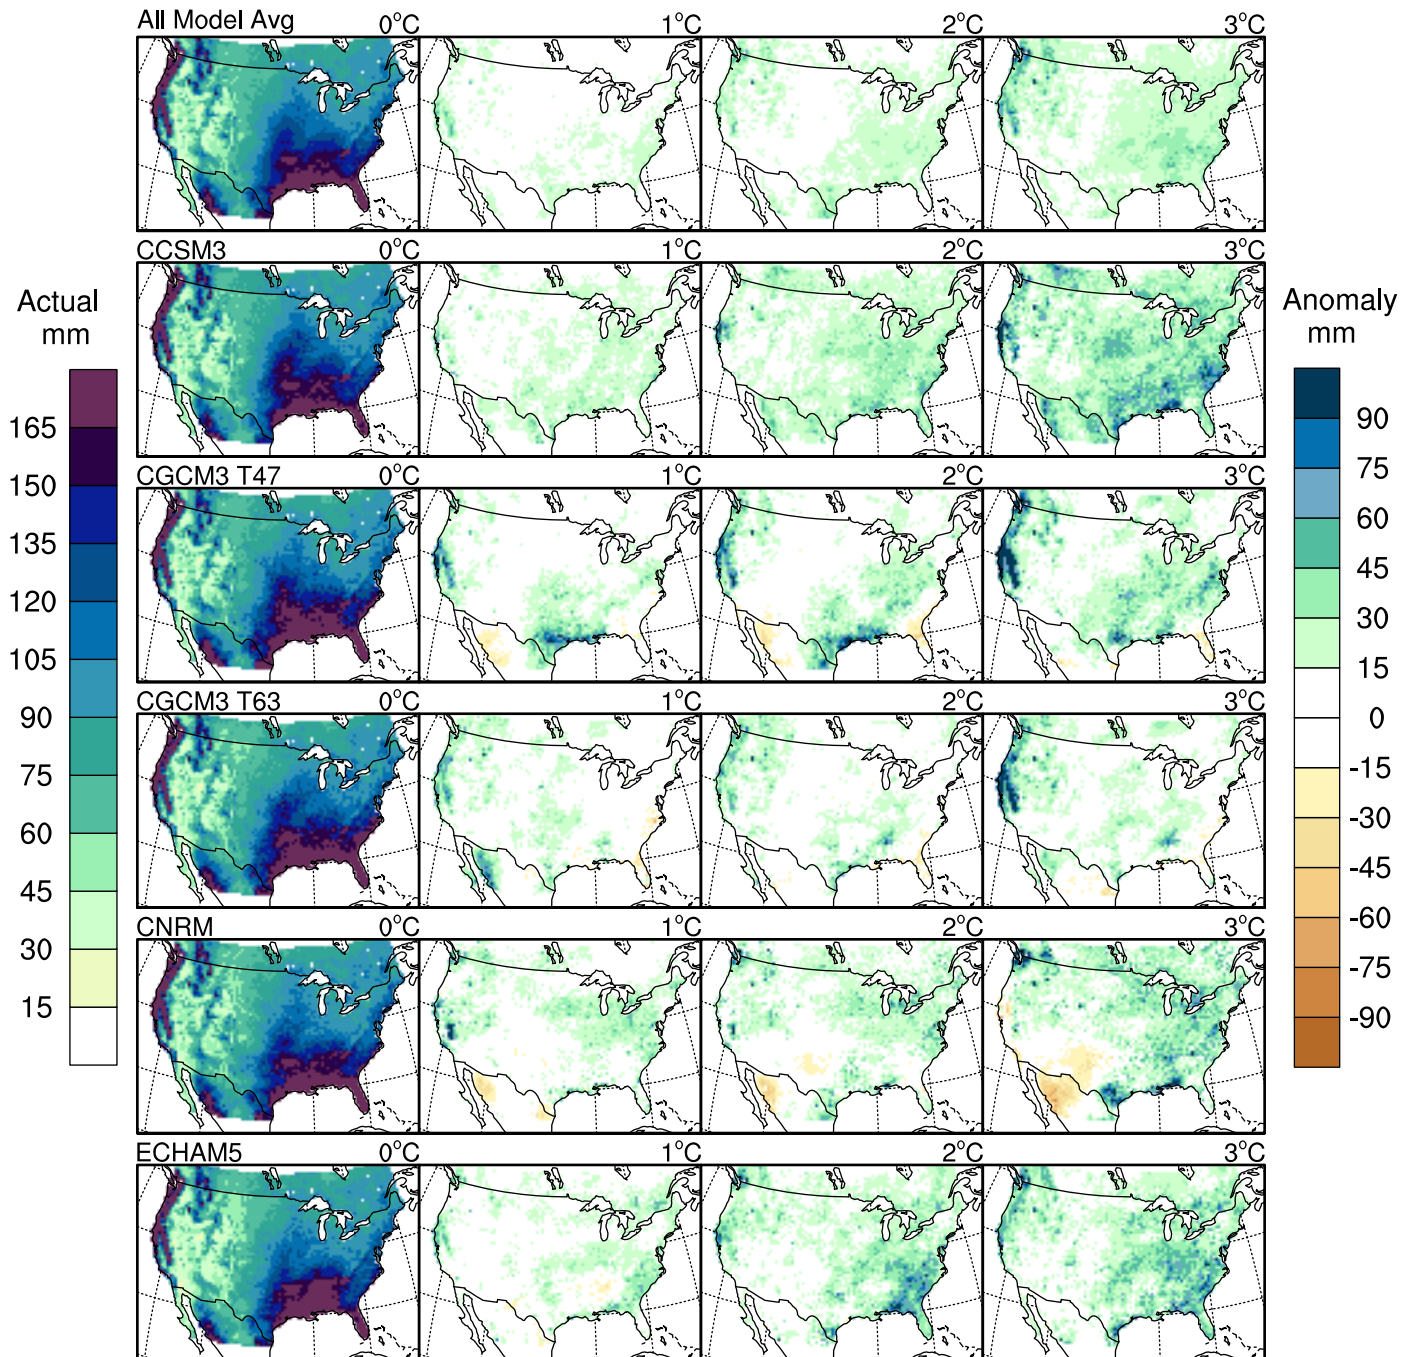
Cumulative Precipitation of the Wettest 5-Day Period



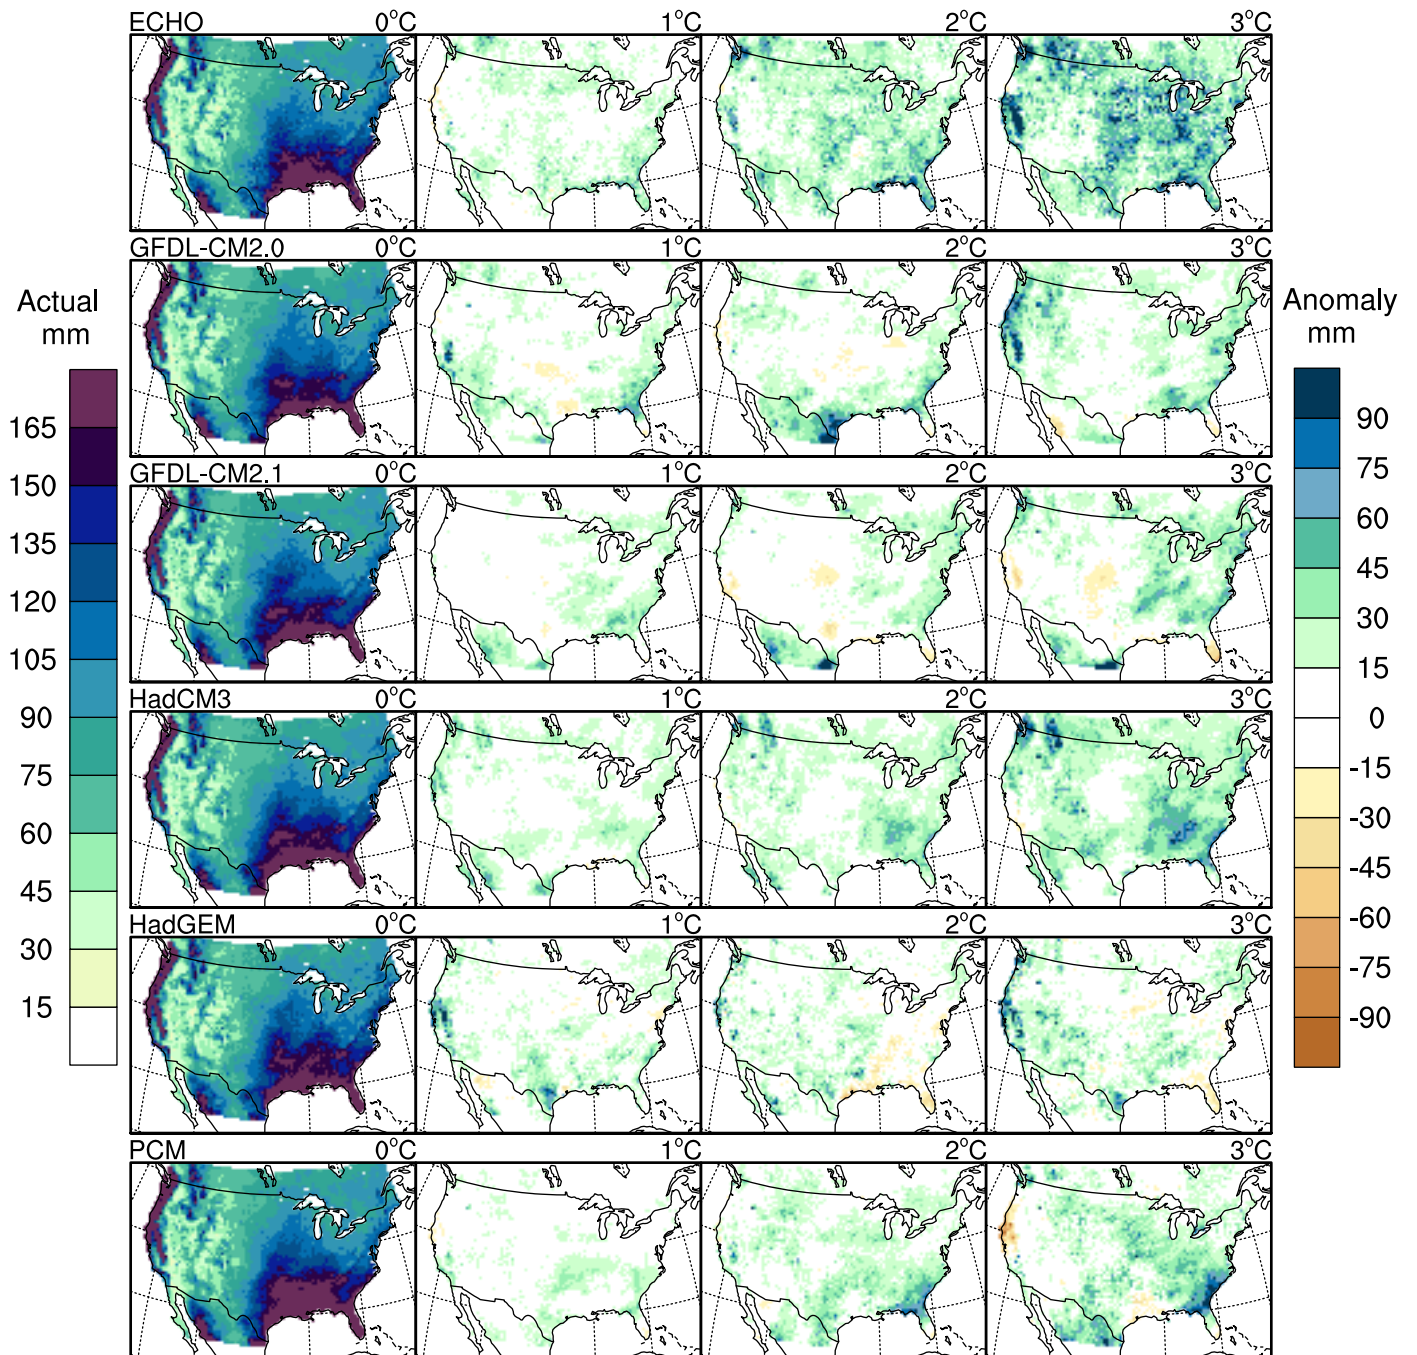
Cumulative Precipitation of the Wettest 5-Day Period



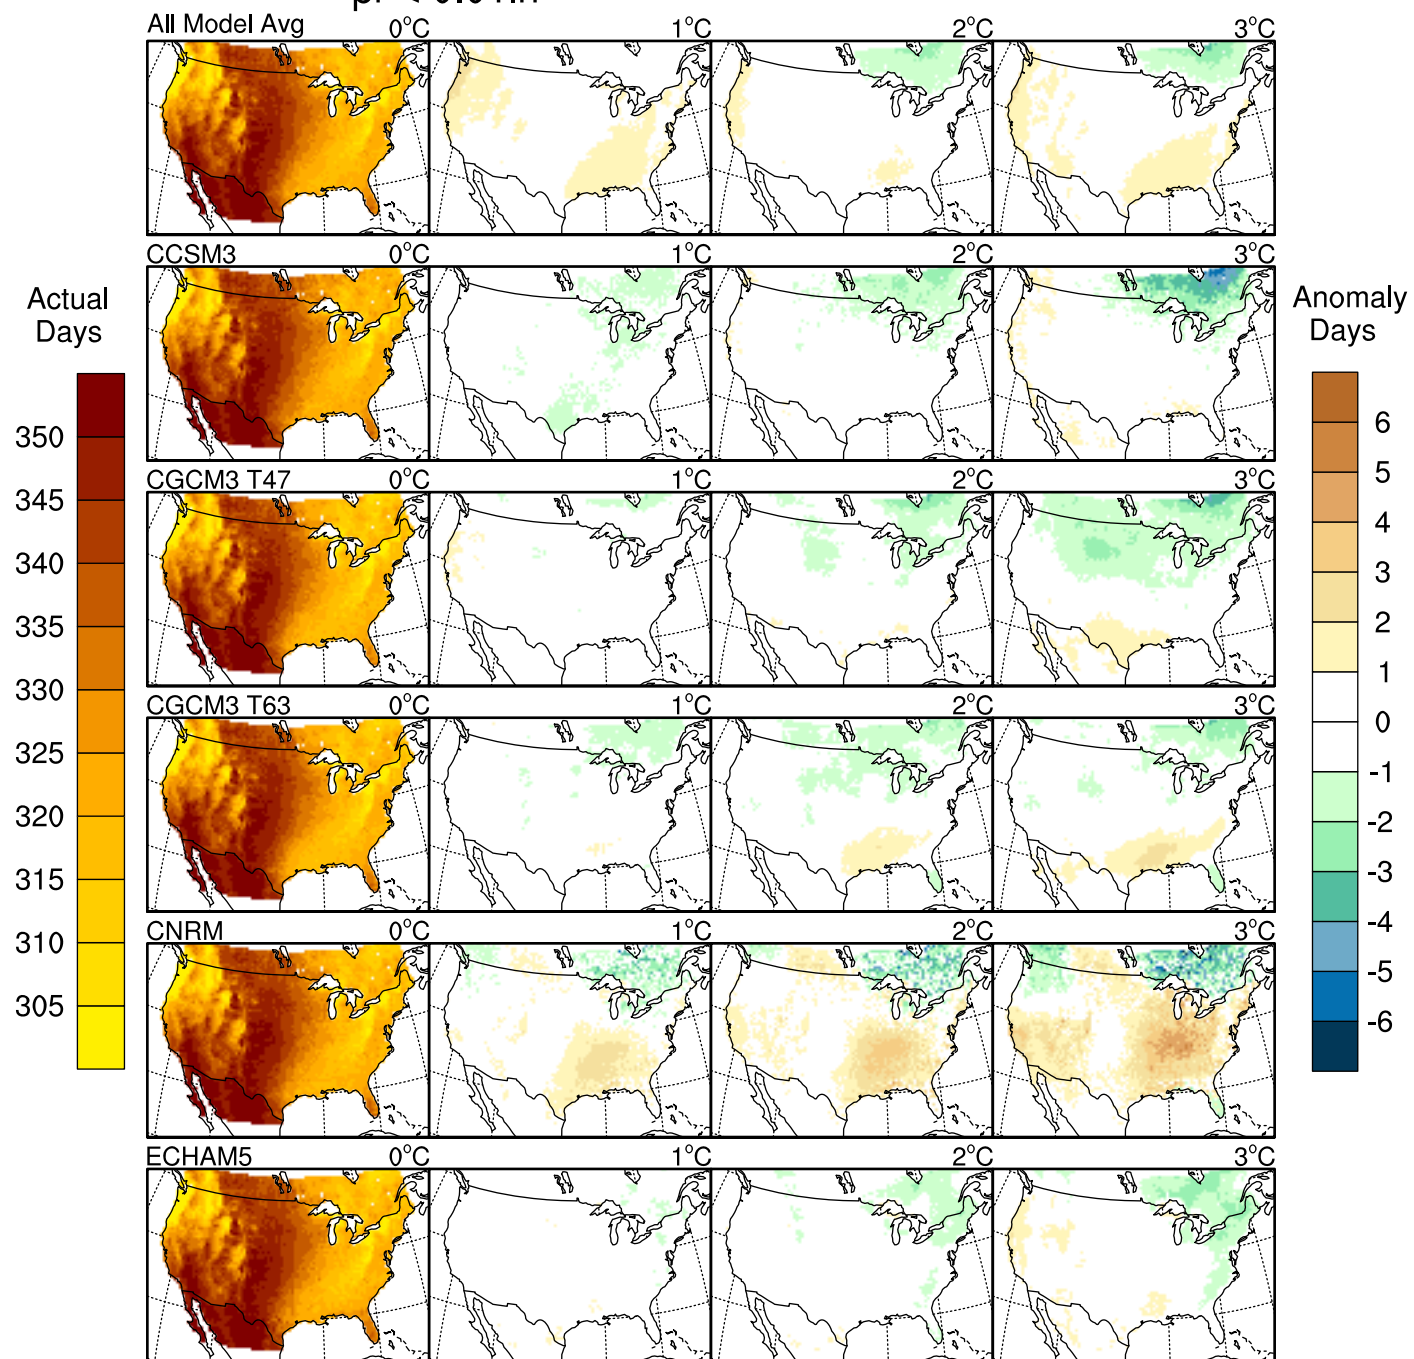
Cumulative Precipitation of the Wettest 10-Day Period



Cumulative Precipitation of the Wettest 10-Day Period

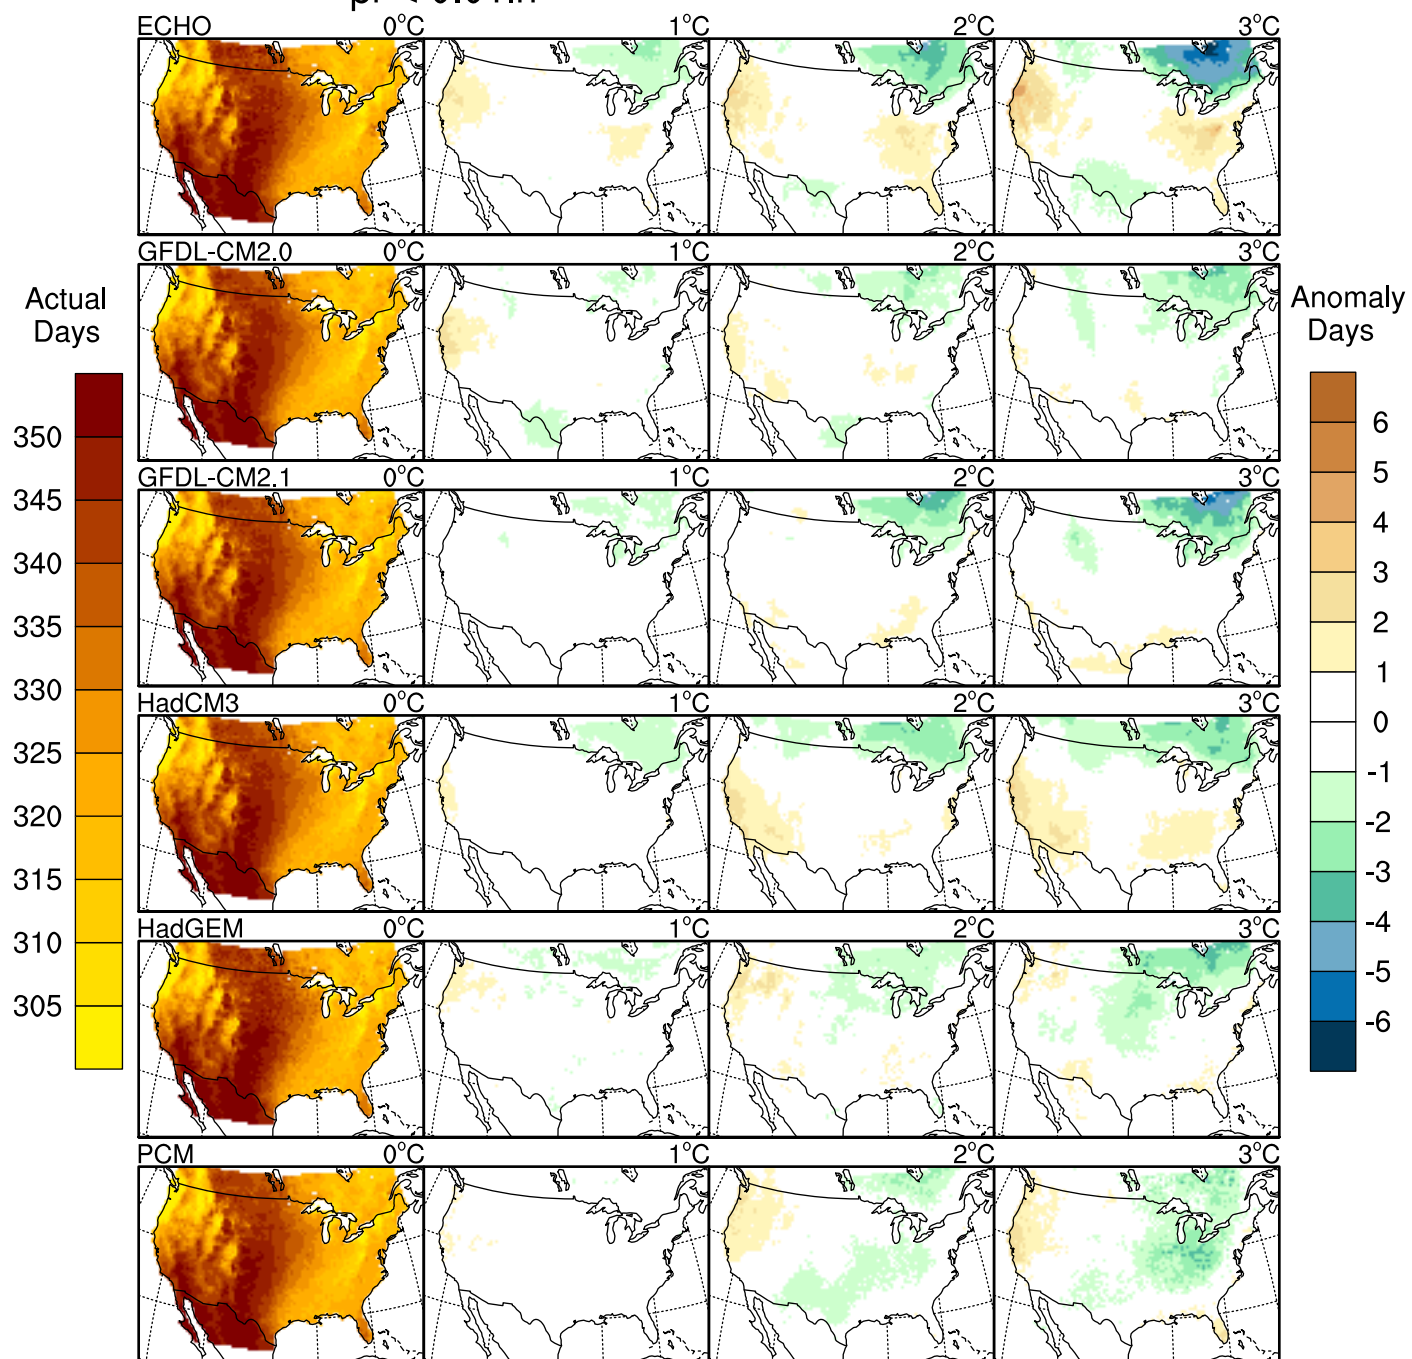


Total Number of Dry Days in Winter (DJF) pr < 0.01in



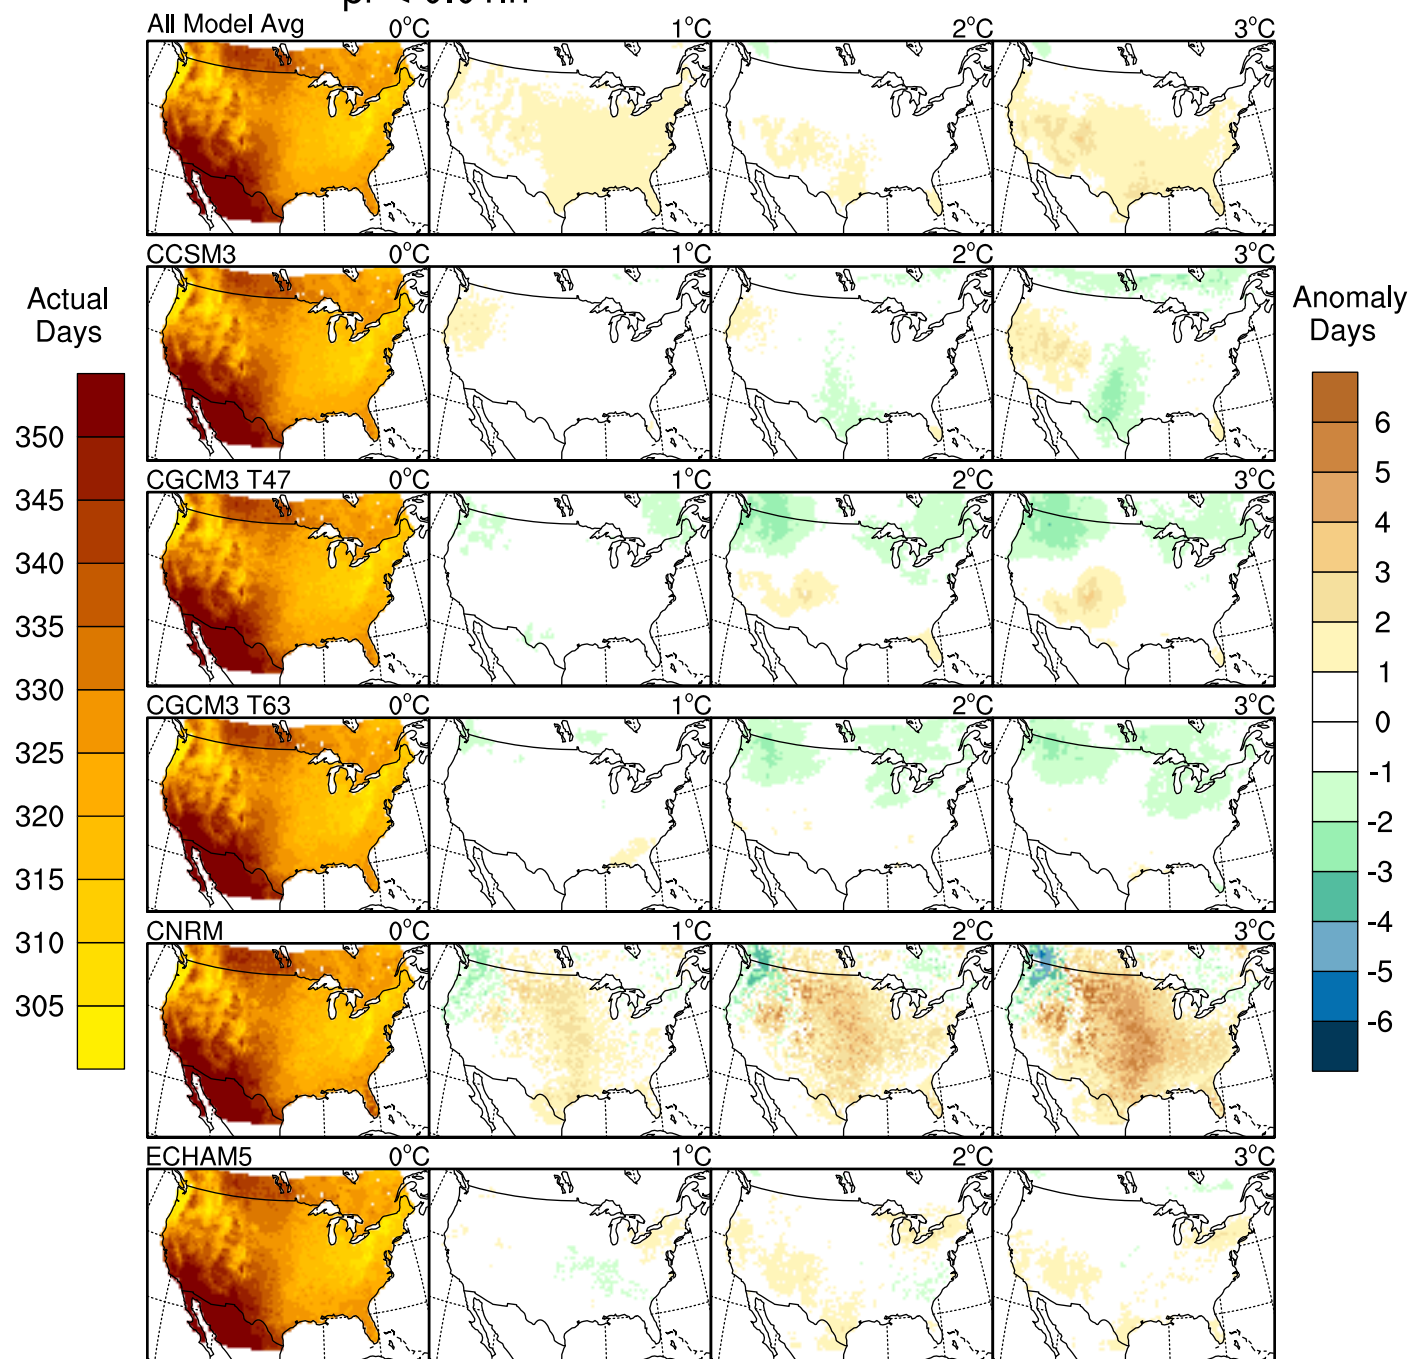
Total Number of Dry Days in Winter (DJF)

$pr < 0.01in$



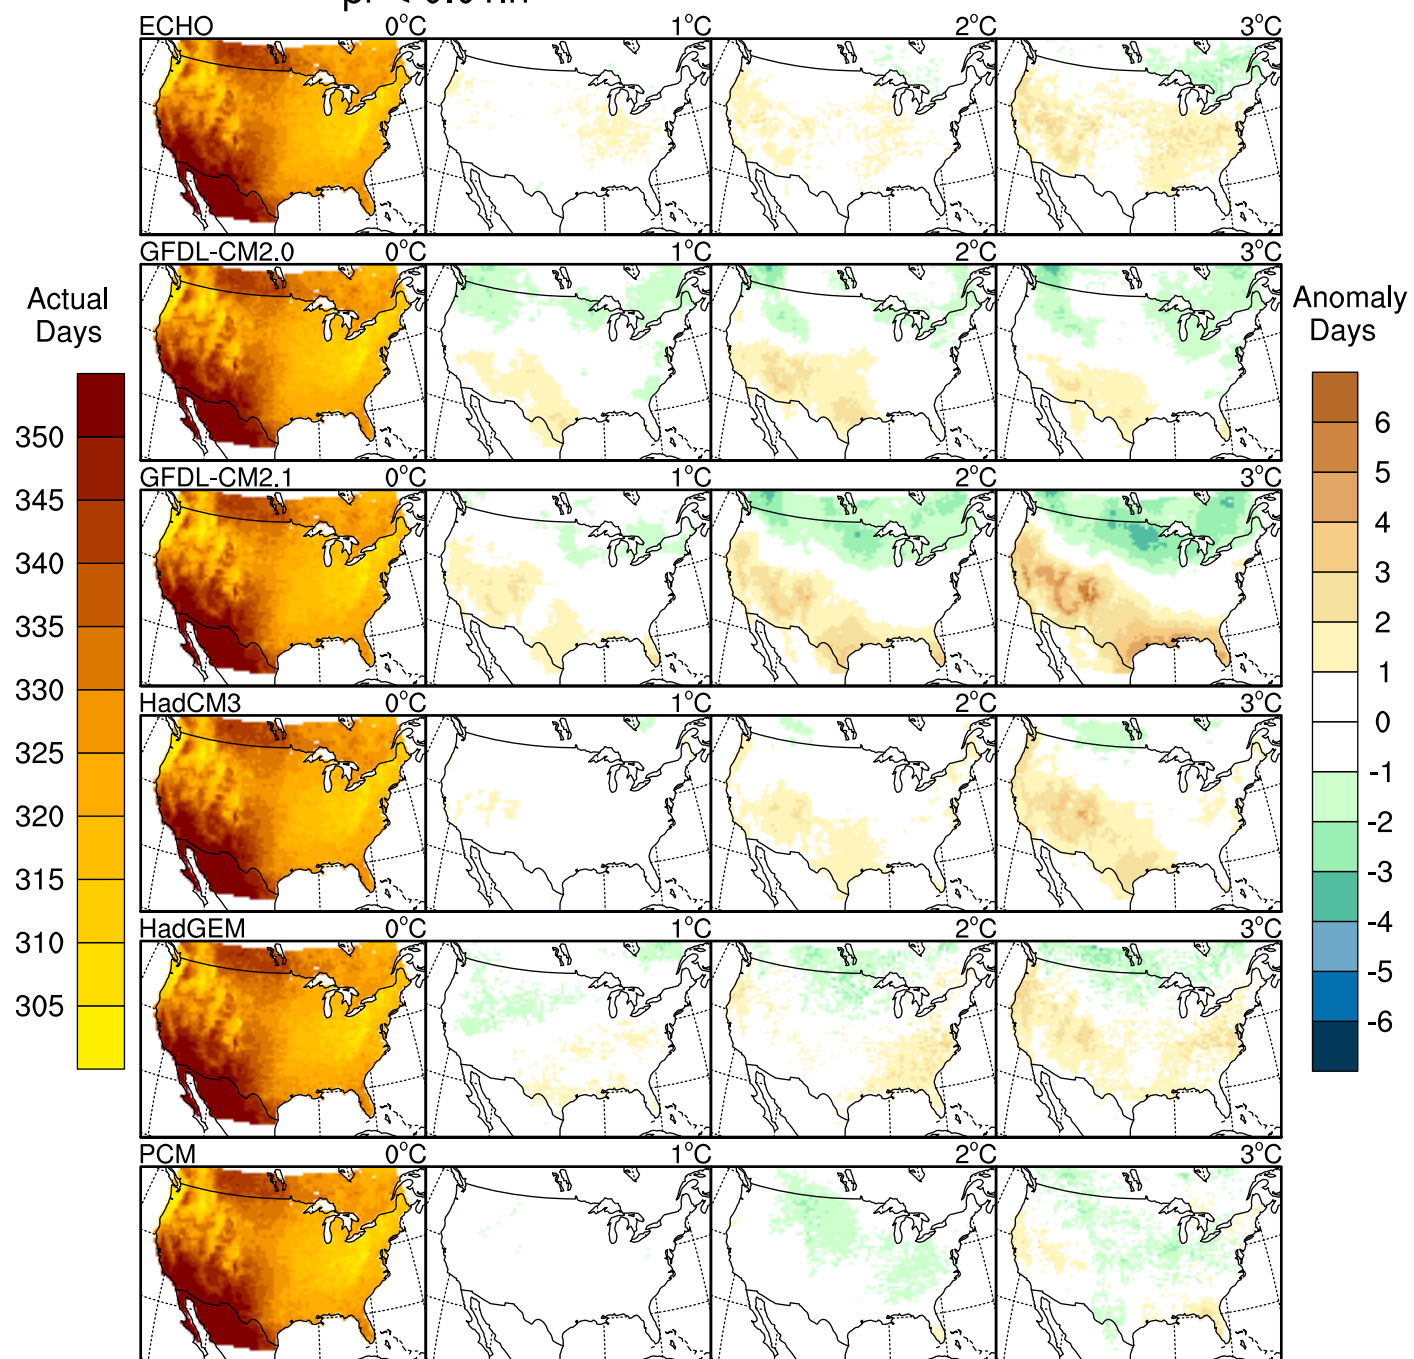
Total Number of Dry Days in Spring (MAM)

pr < 0.01in



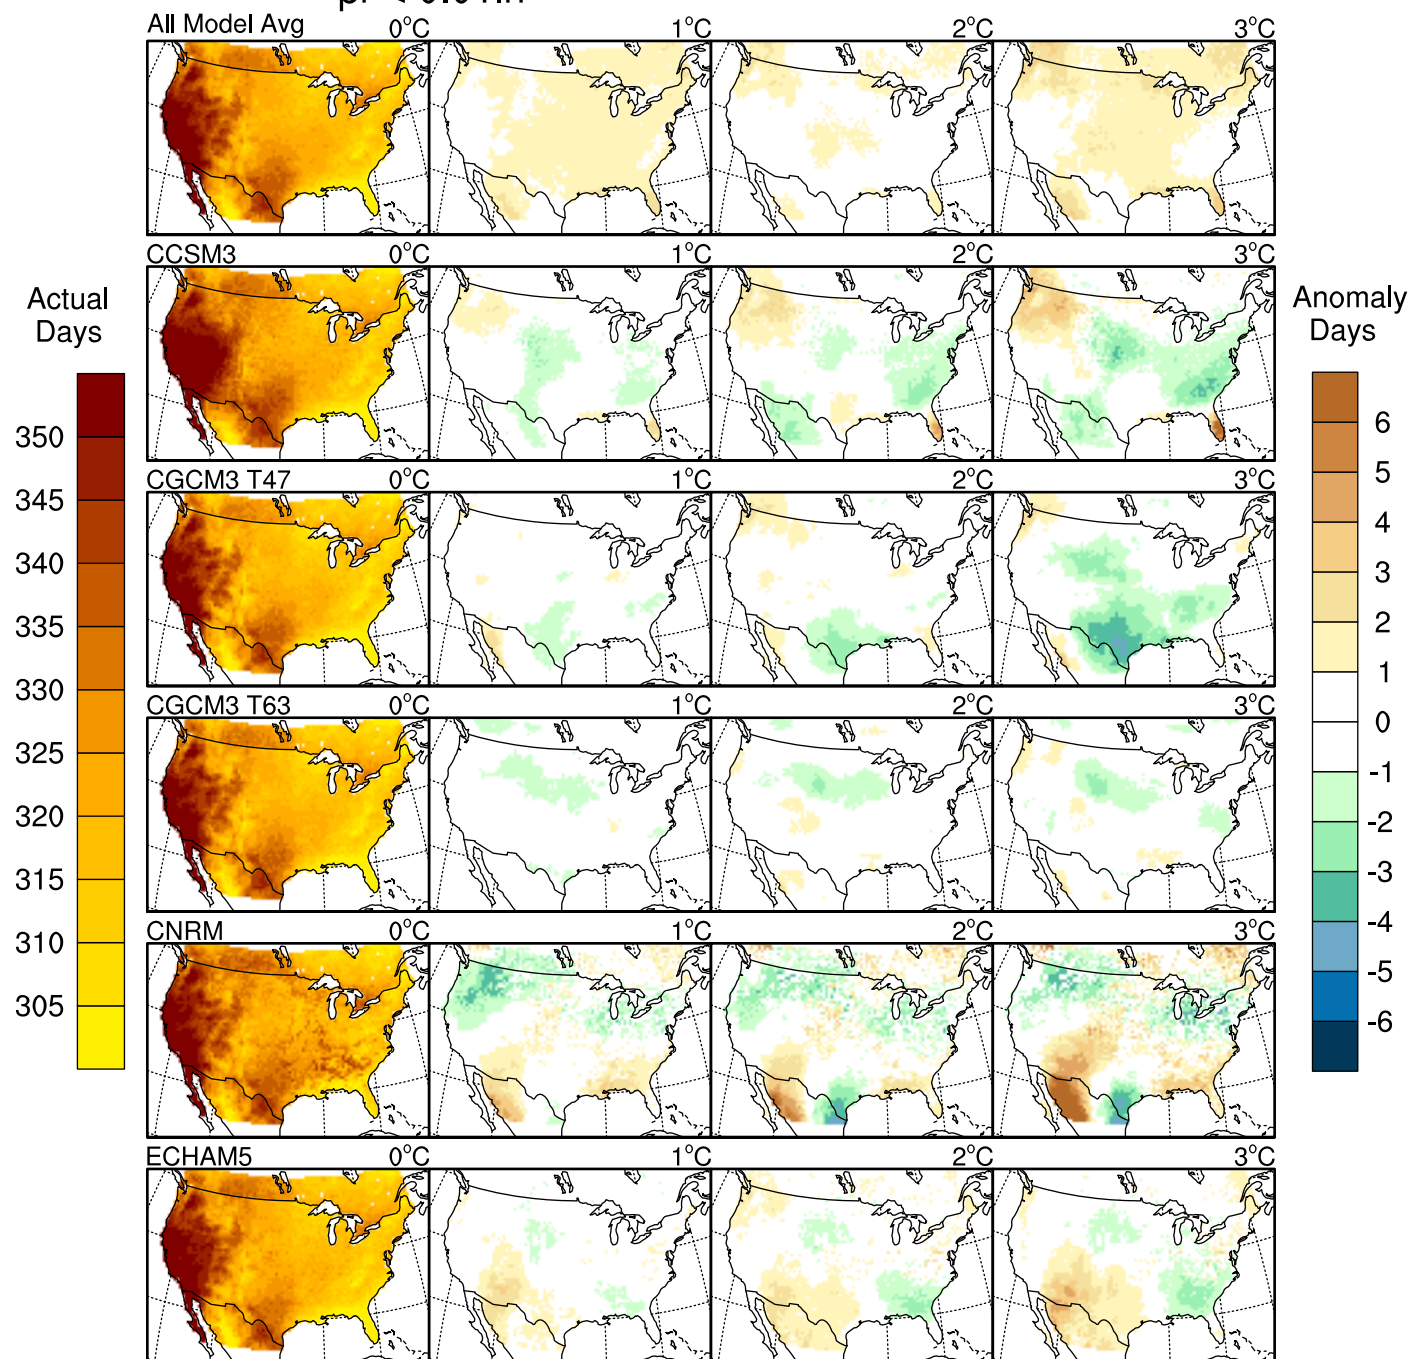
Total Number of Dry Days in Spring (MAM)

pr < 0.01in



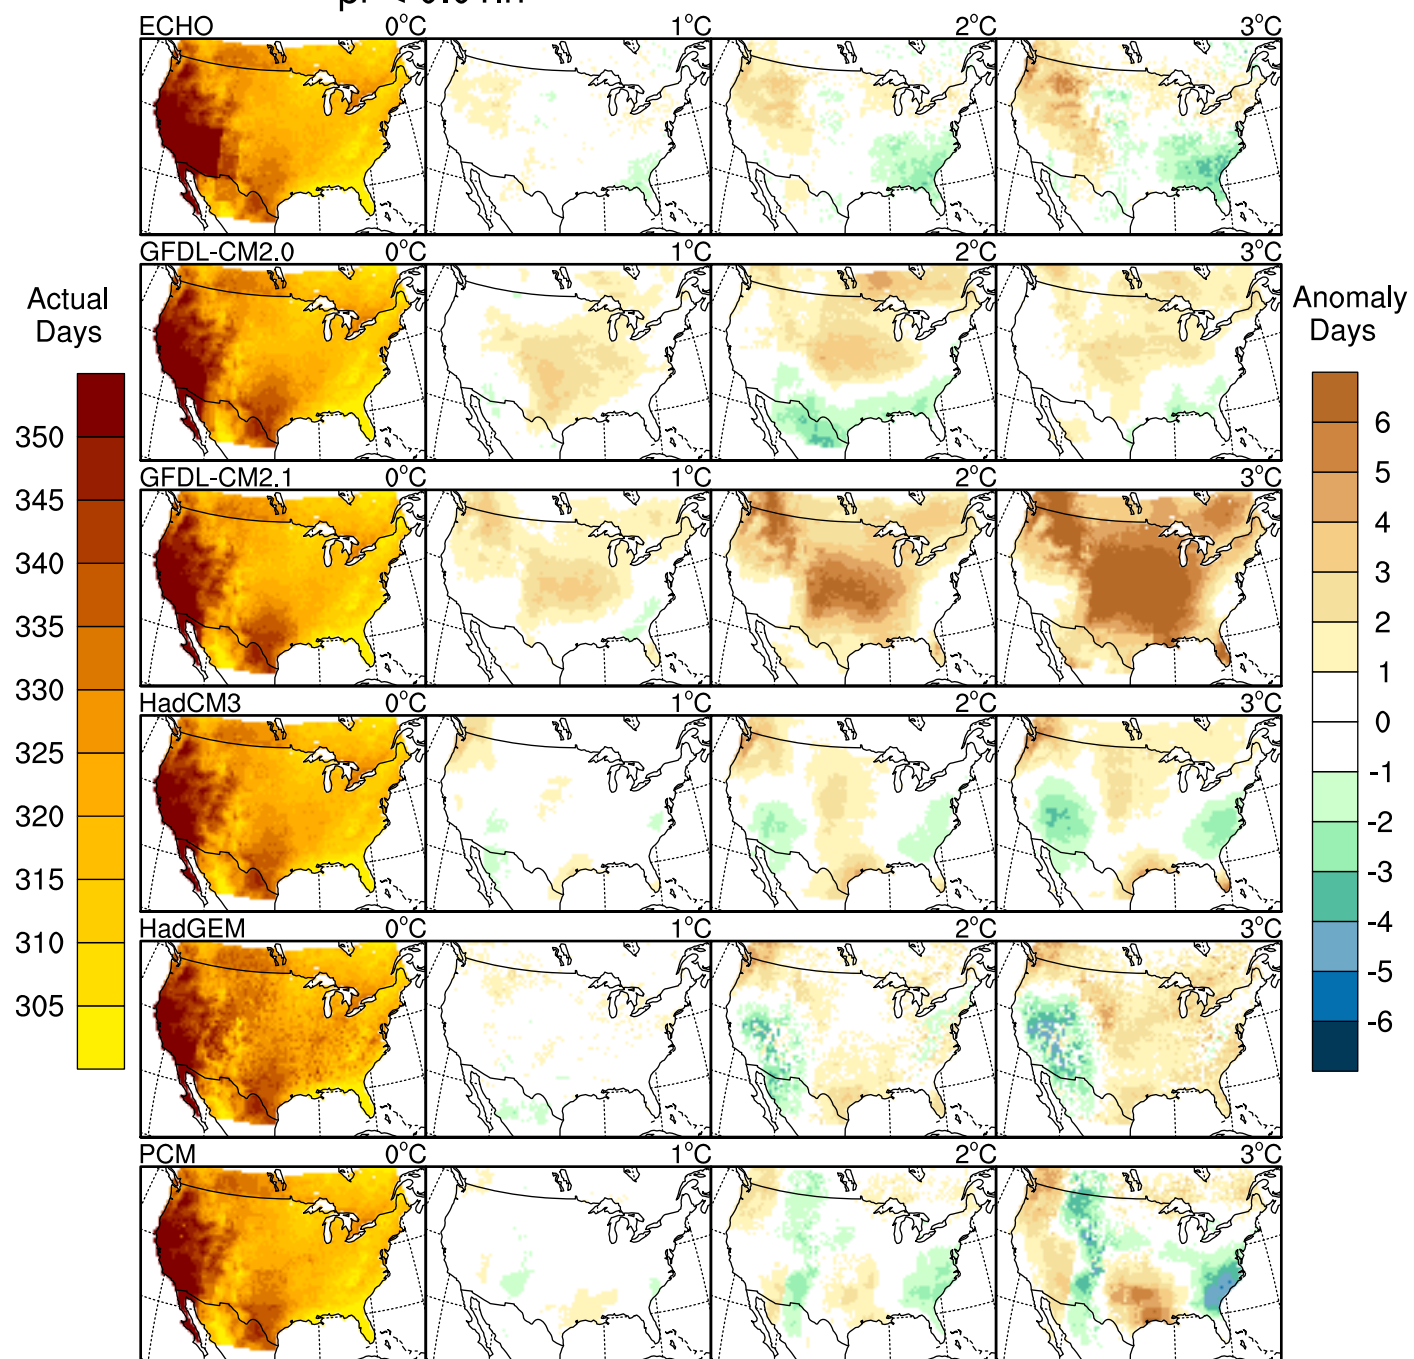
Total Number of Dry Days in Summer (JJA)

pr < 0.01in



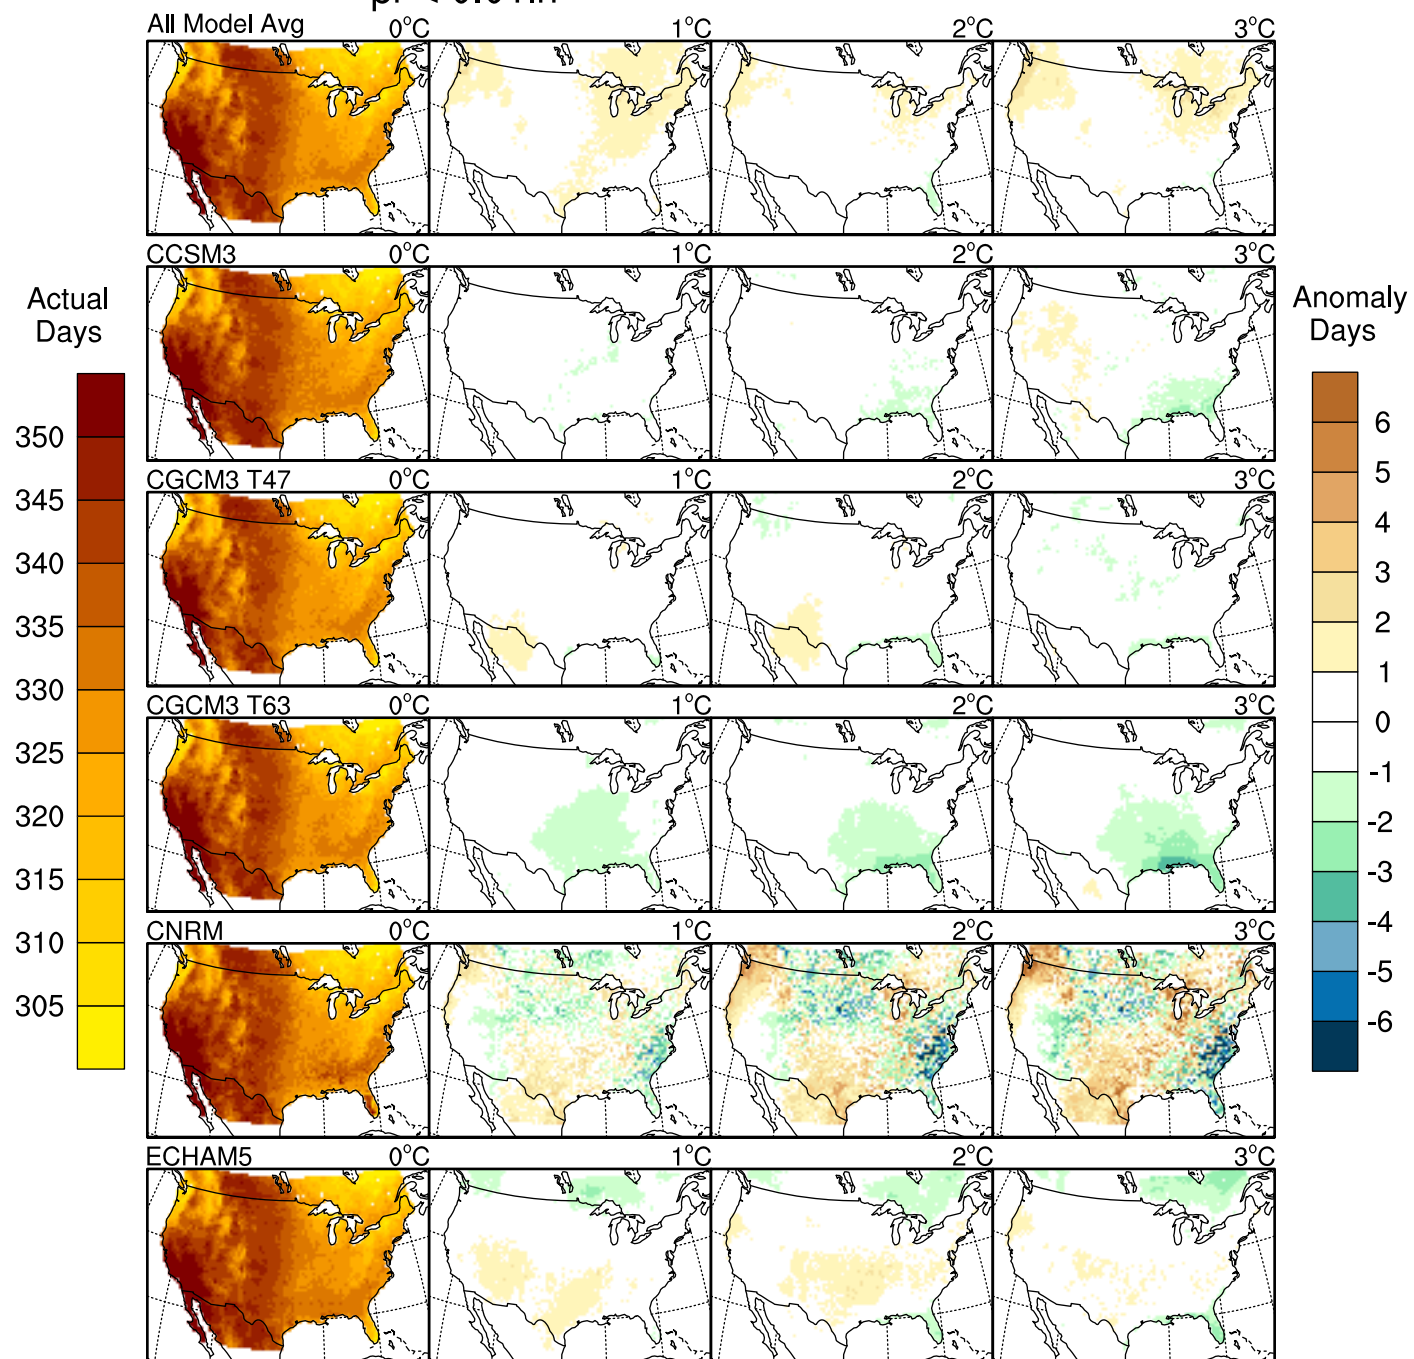
Total Number of Dry Days in Summer (JJA)

pr < 0.01in



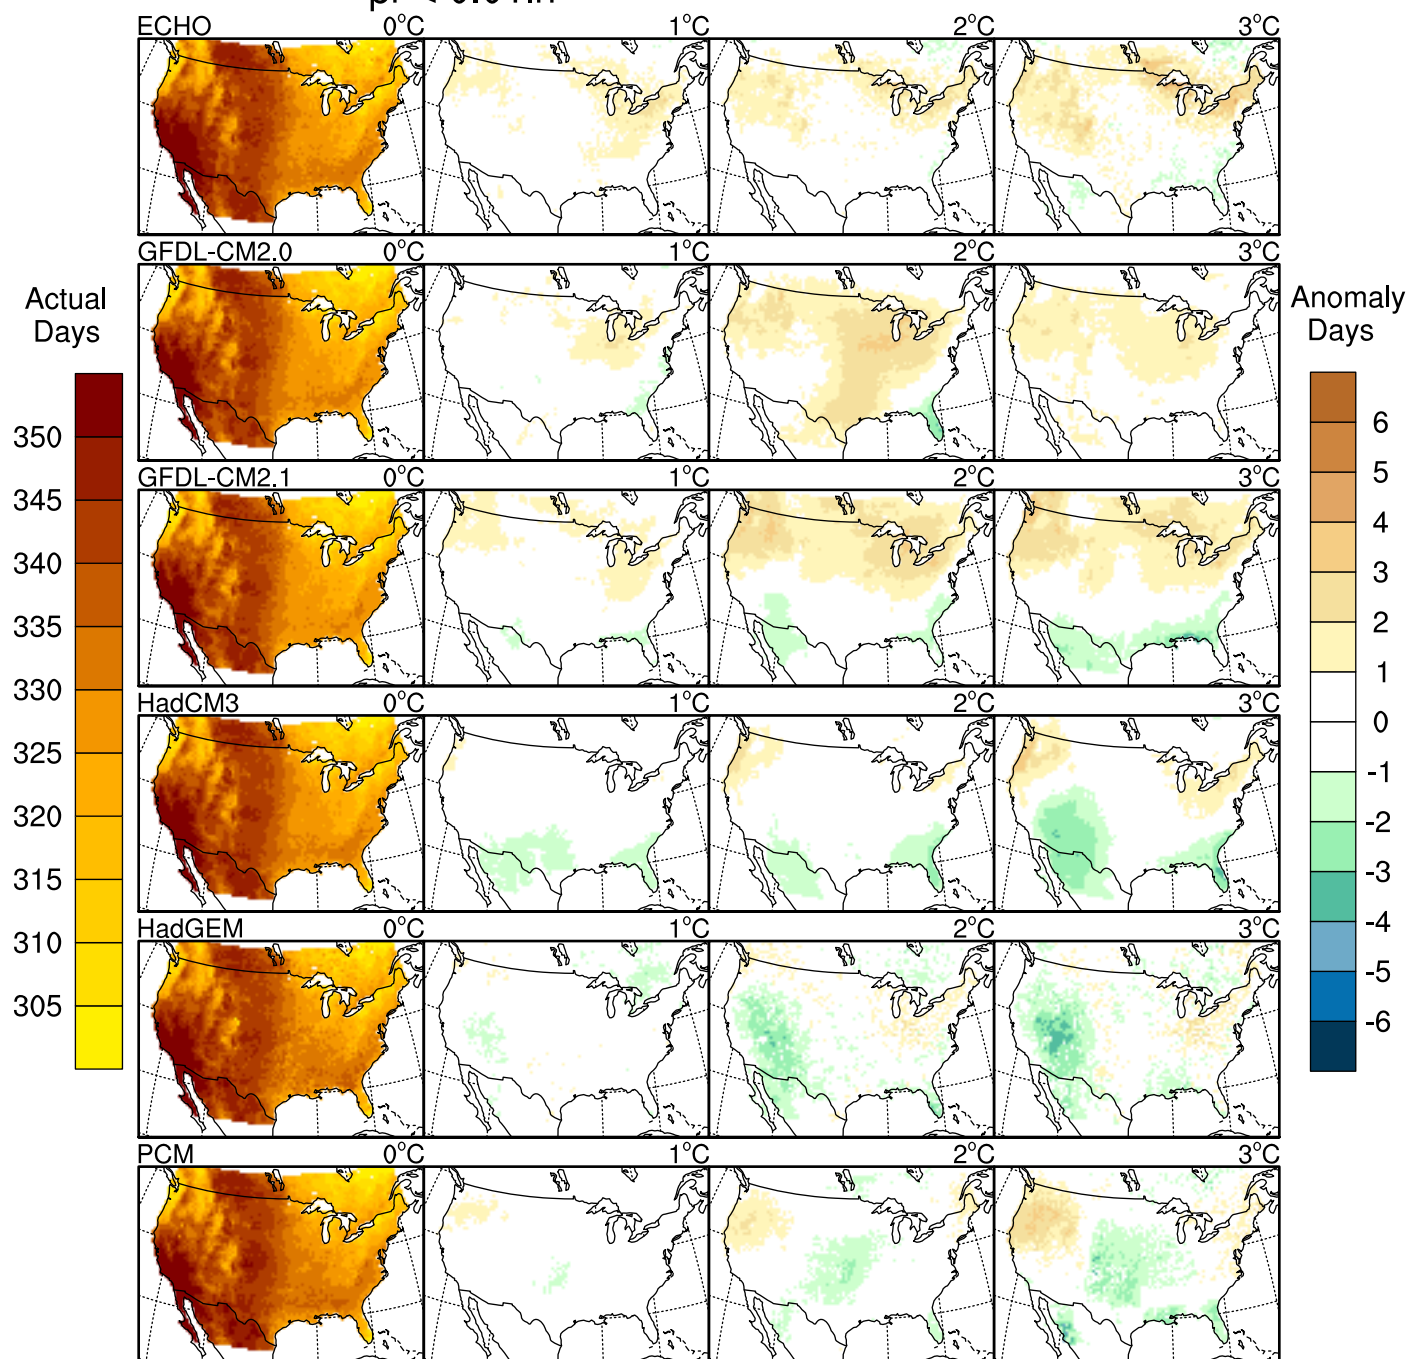
Total Number of Dry Days in Fall (SON)

pr < 0.01in

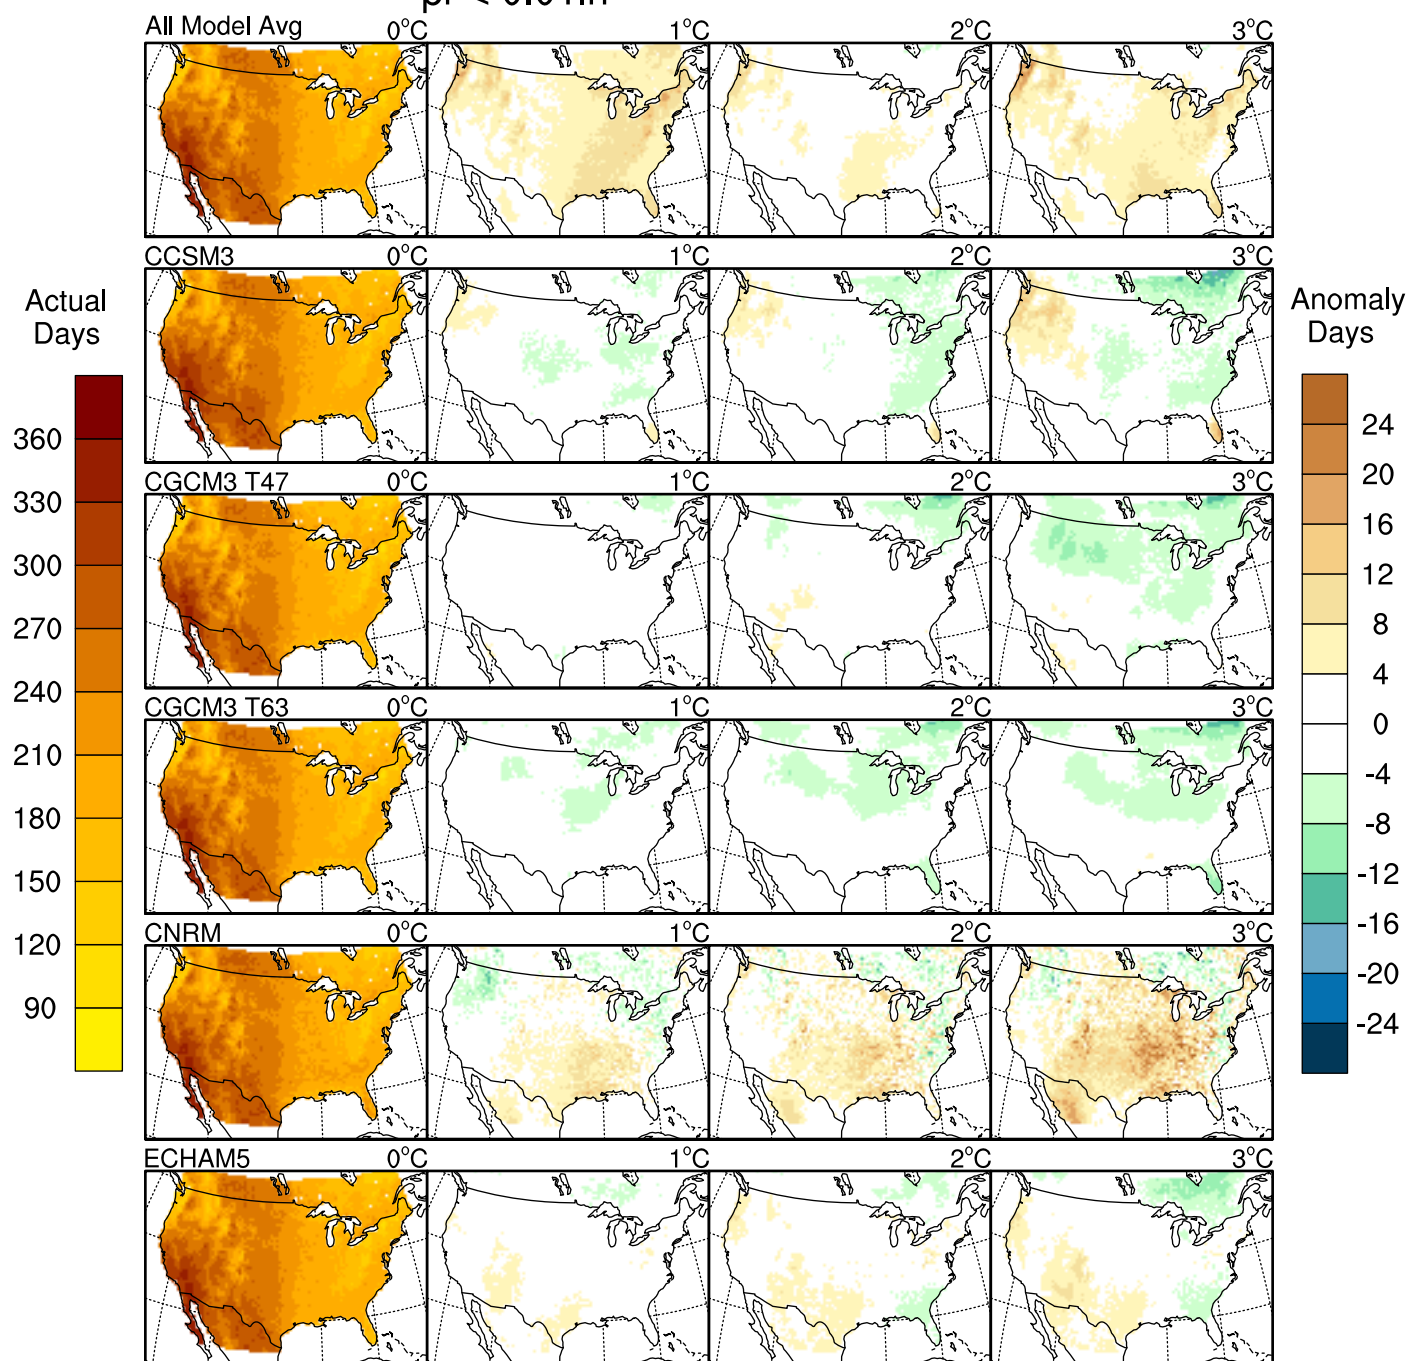


Total Number of Dry Days in Fall (SON)

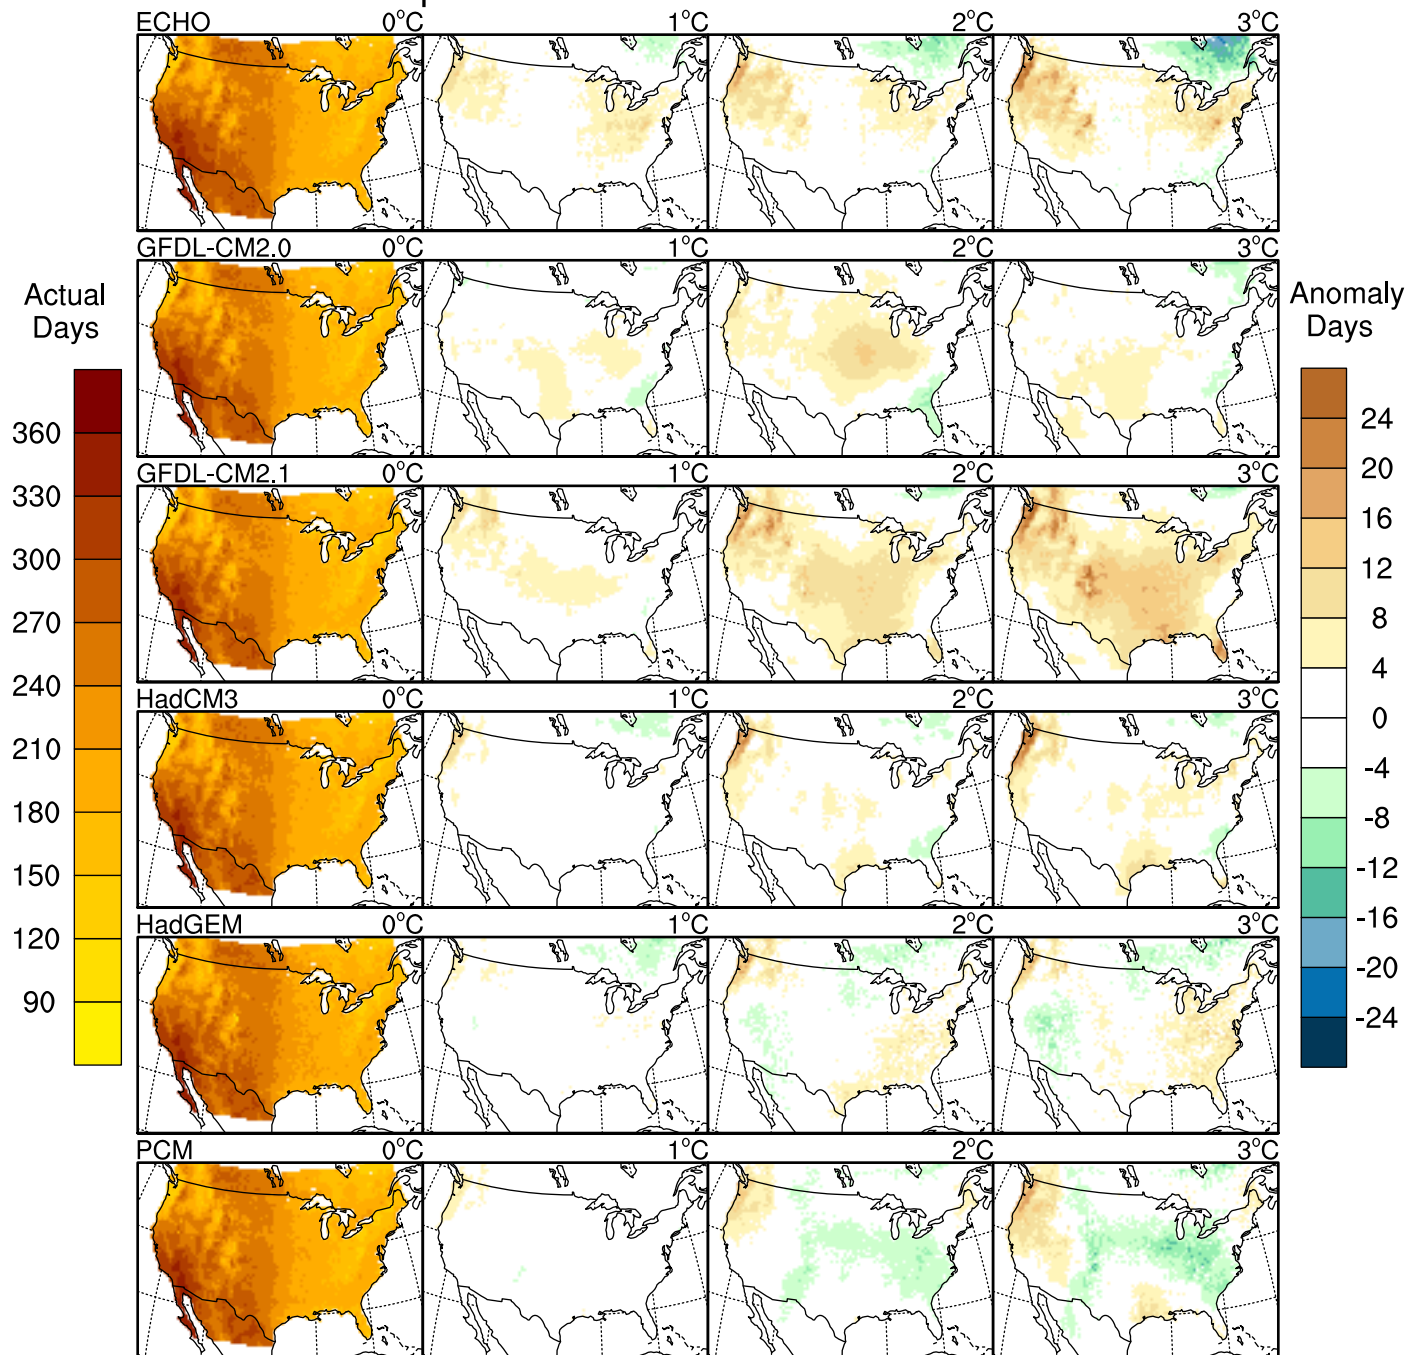
pr < 0.01in



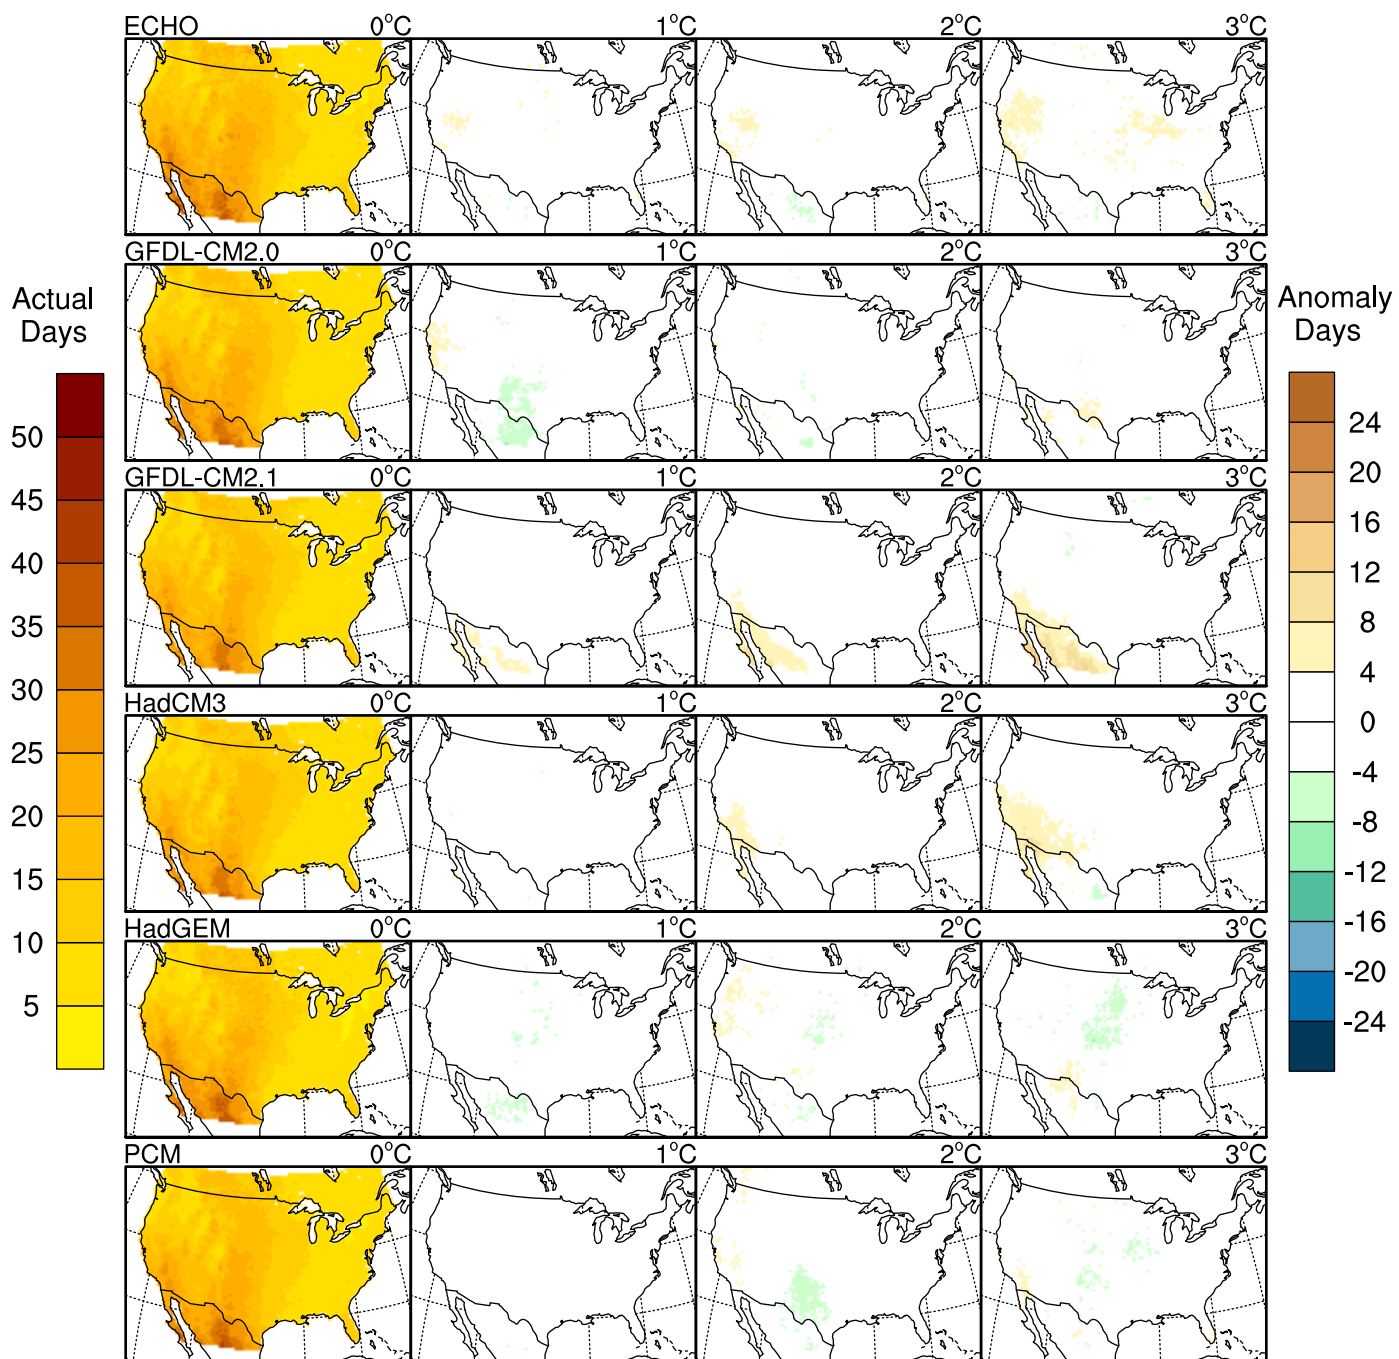
Total Number of Annual Dry Days pr < 0.01in



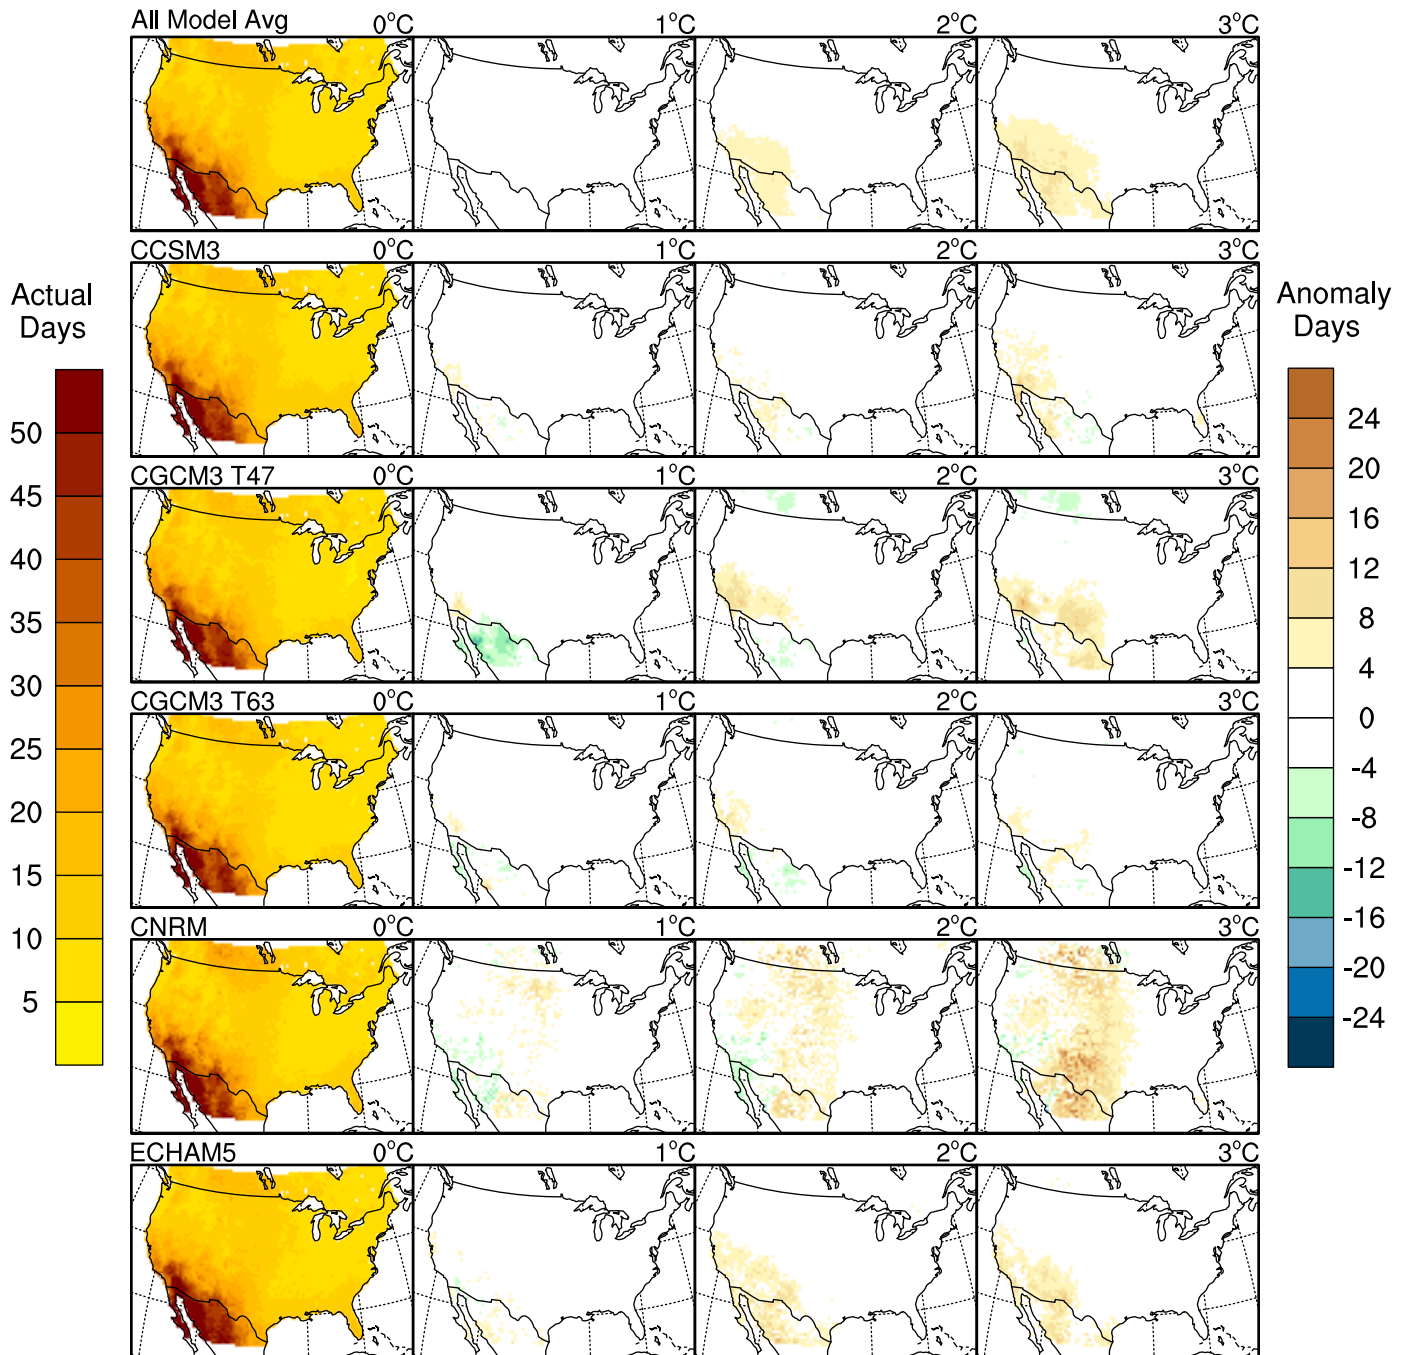
Total Number of Annual Dry Days pr < 0.01in



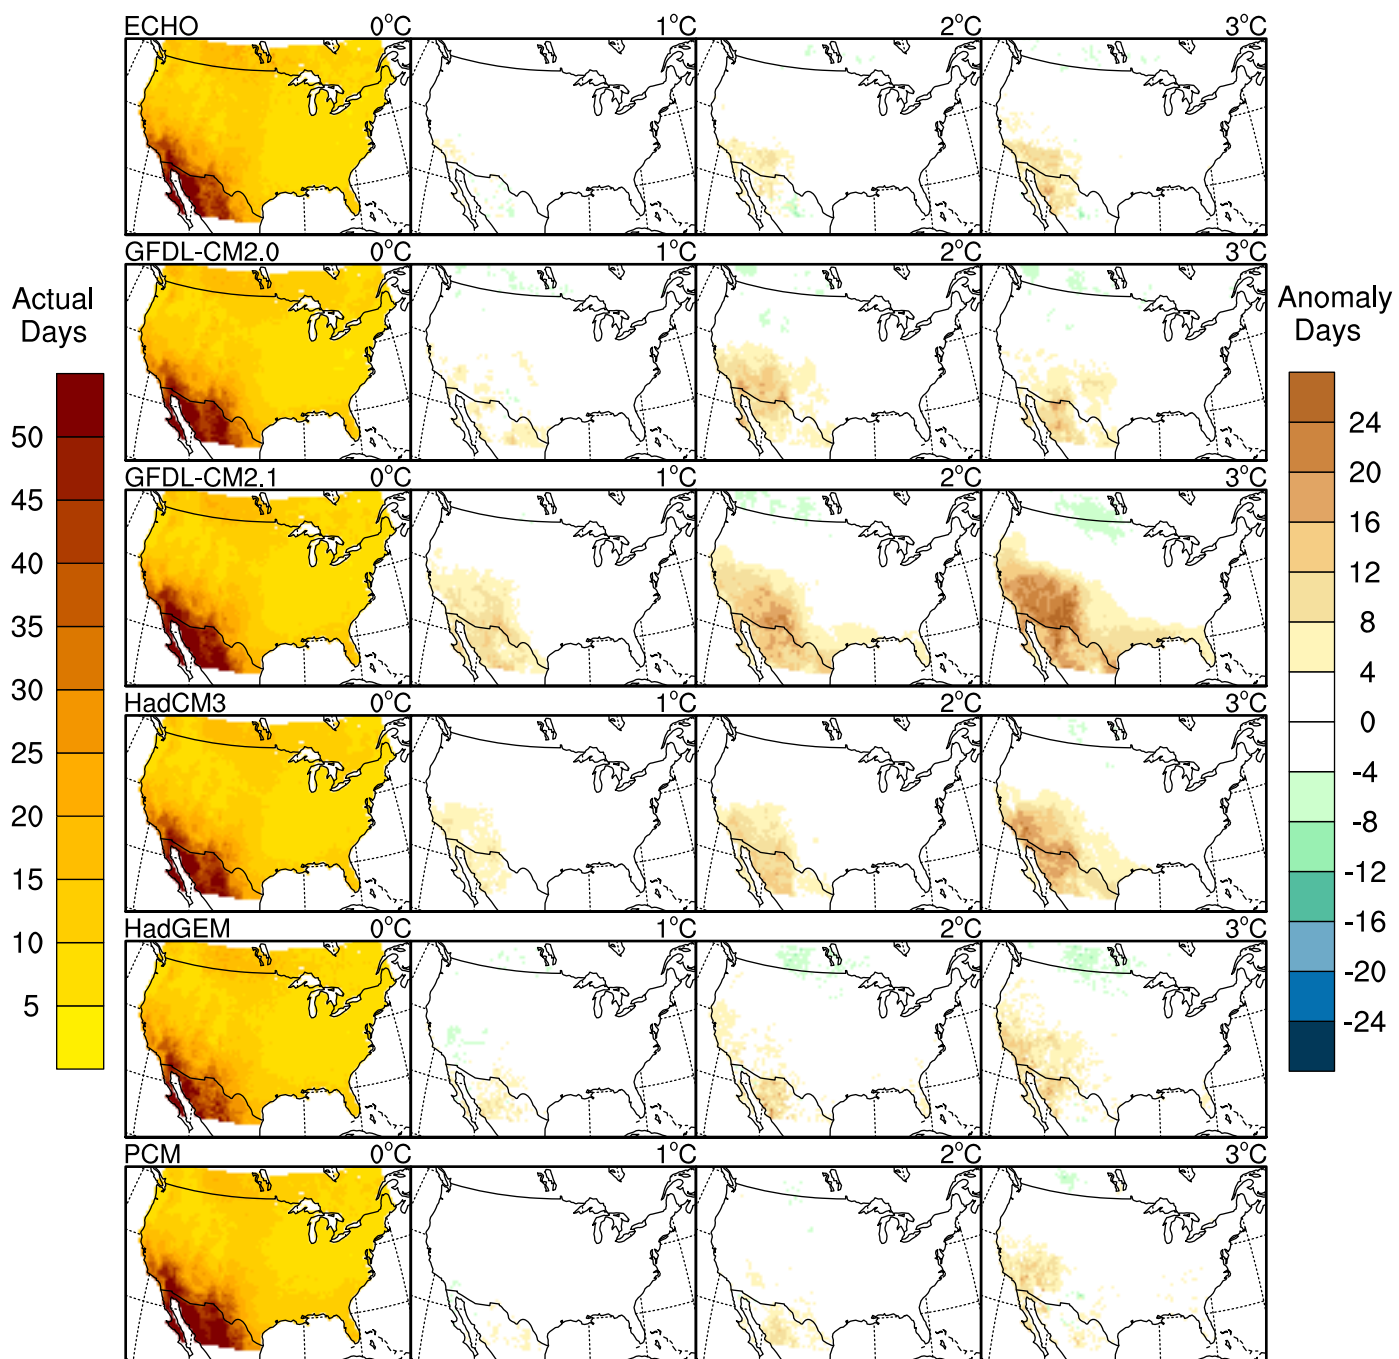
Longest Period of Dry Days < 0.01in (Winter)



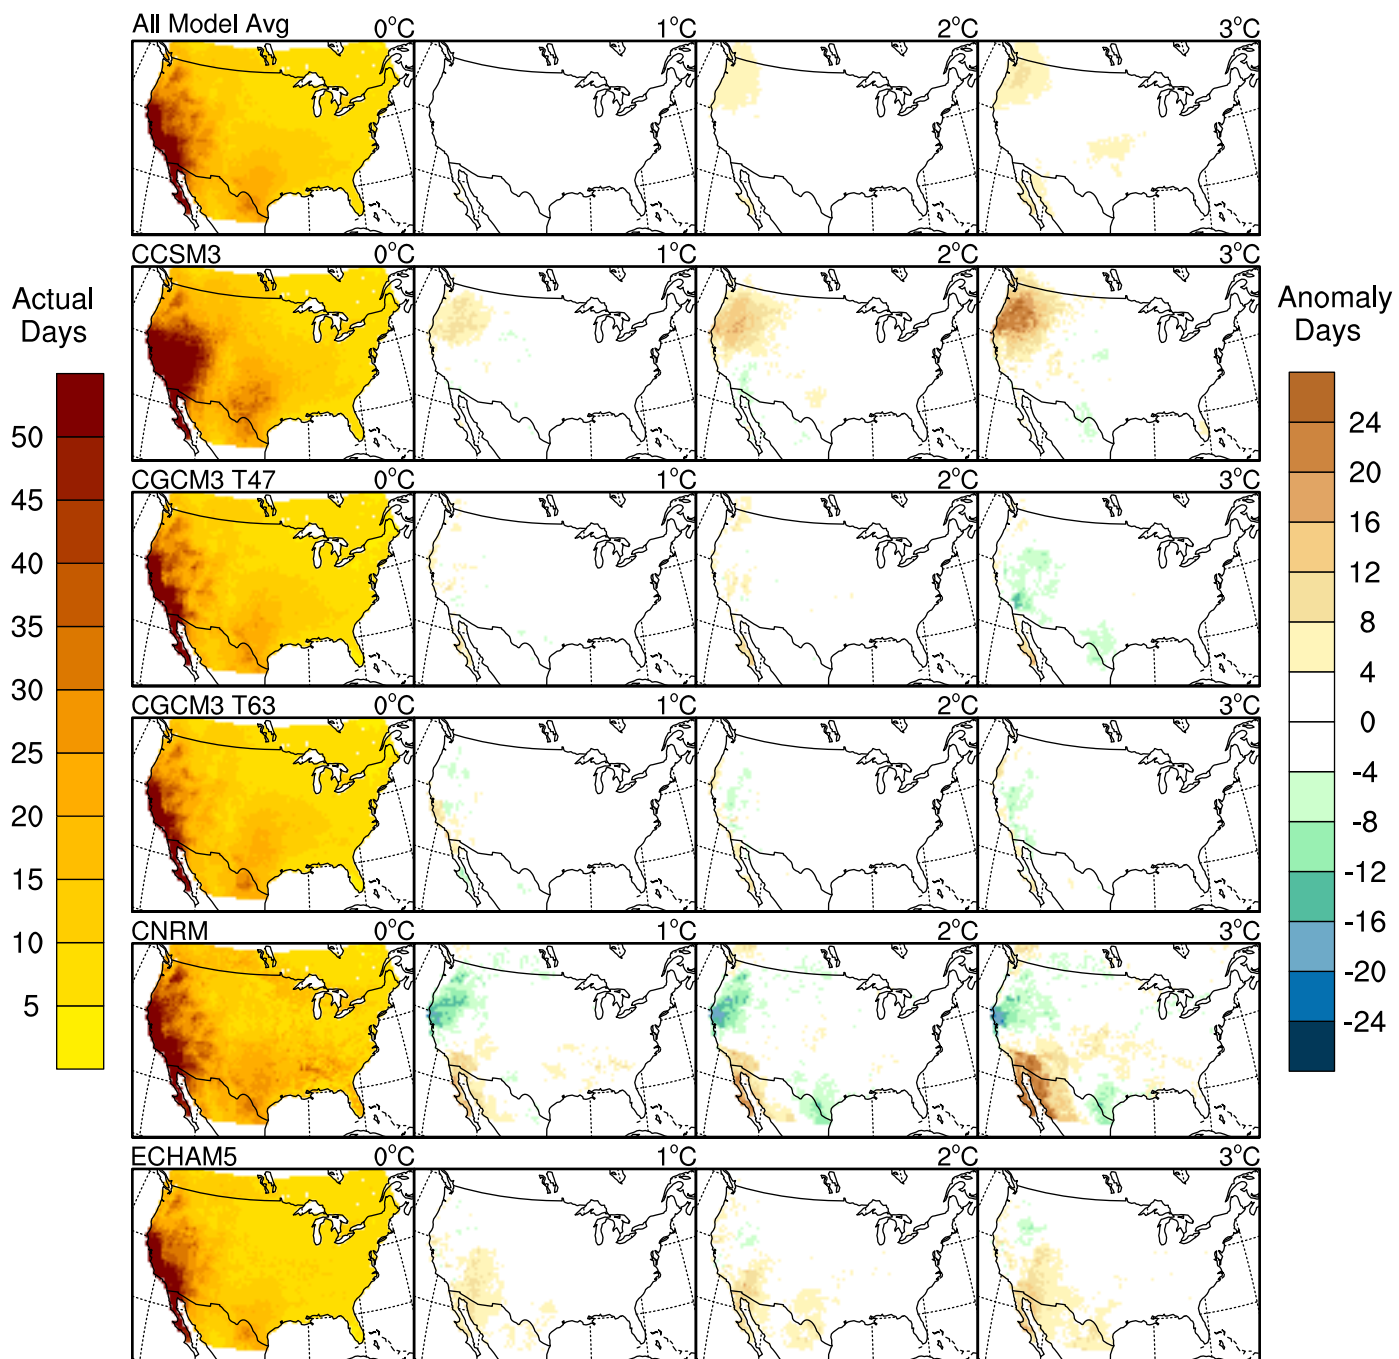
Longest Period of Dry Days < 0.01in (Spring)



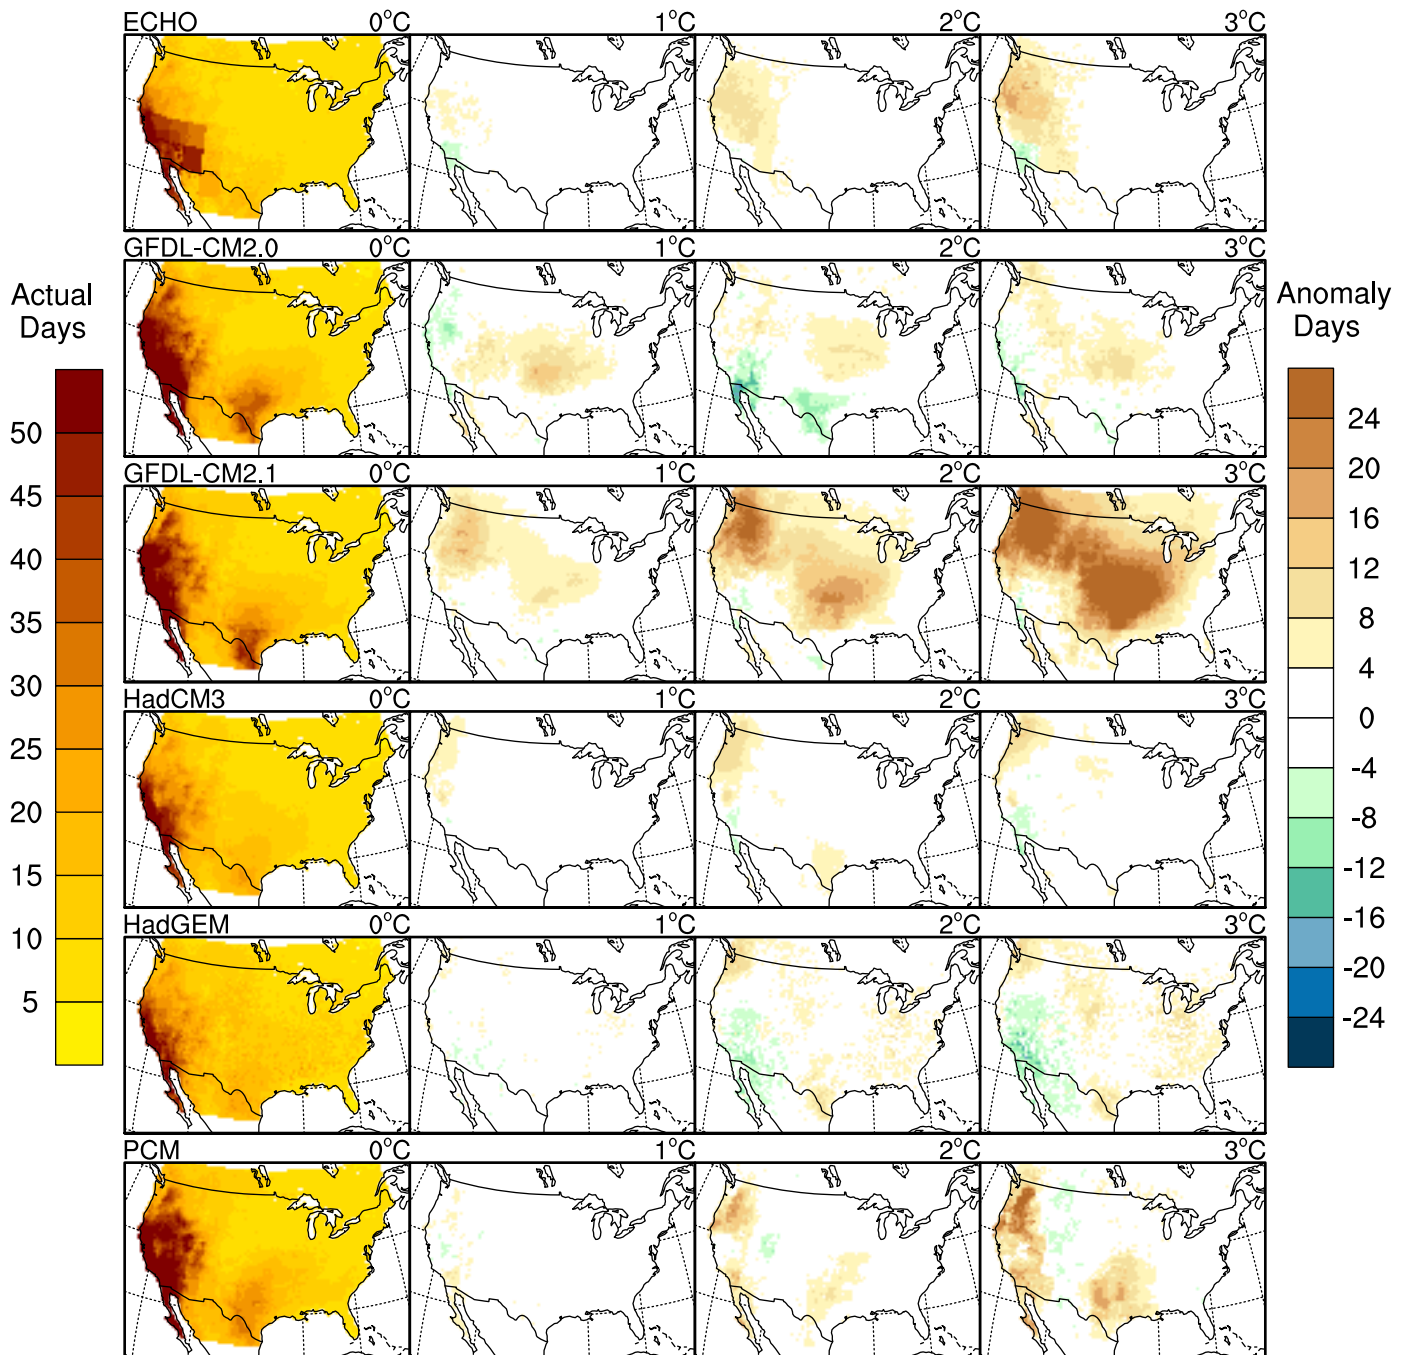
Longest Period of Dry Days < 0.01in (Spring)



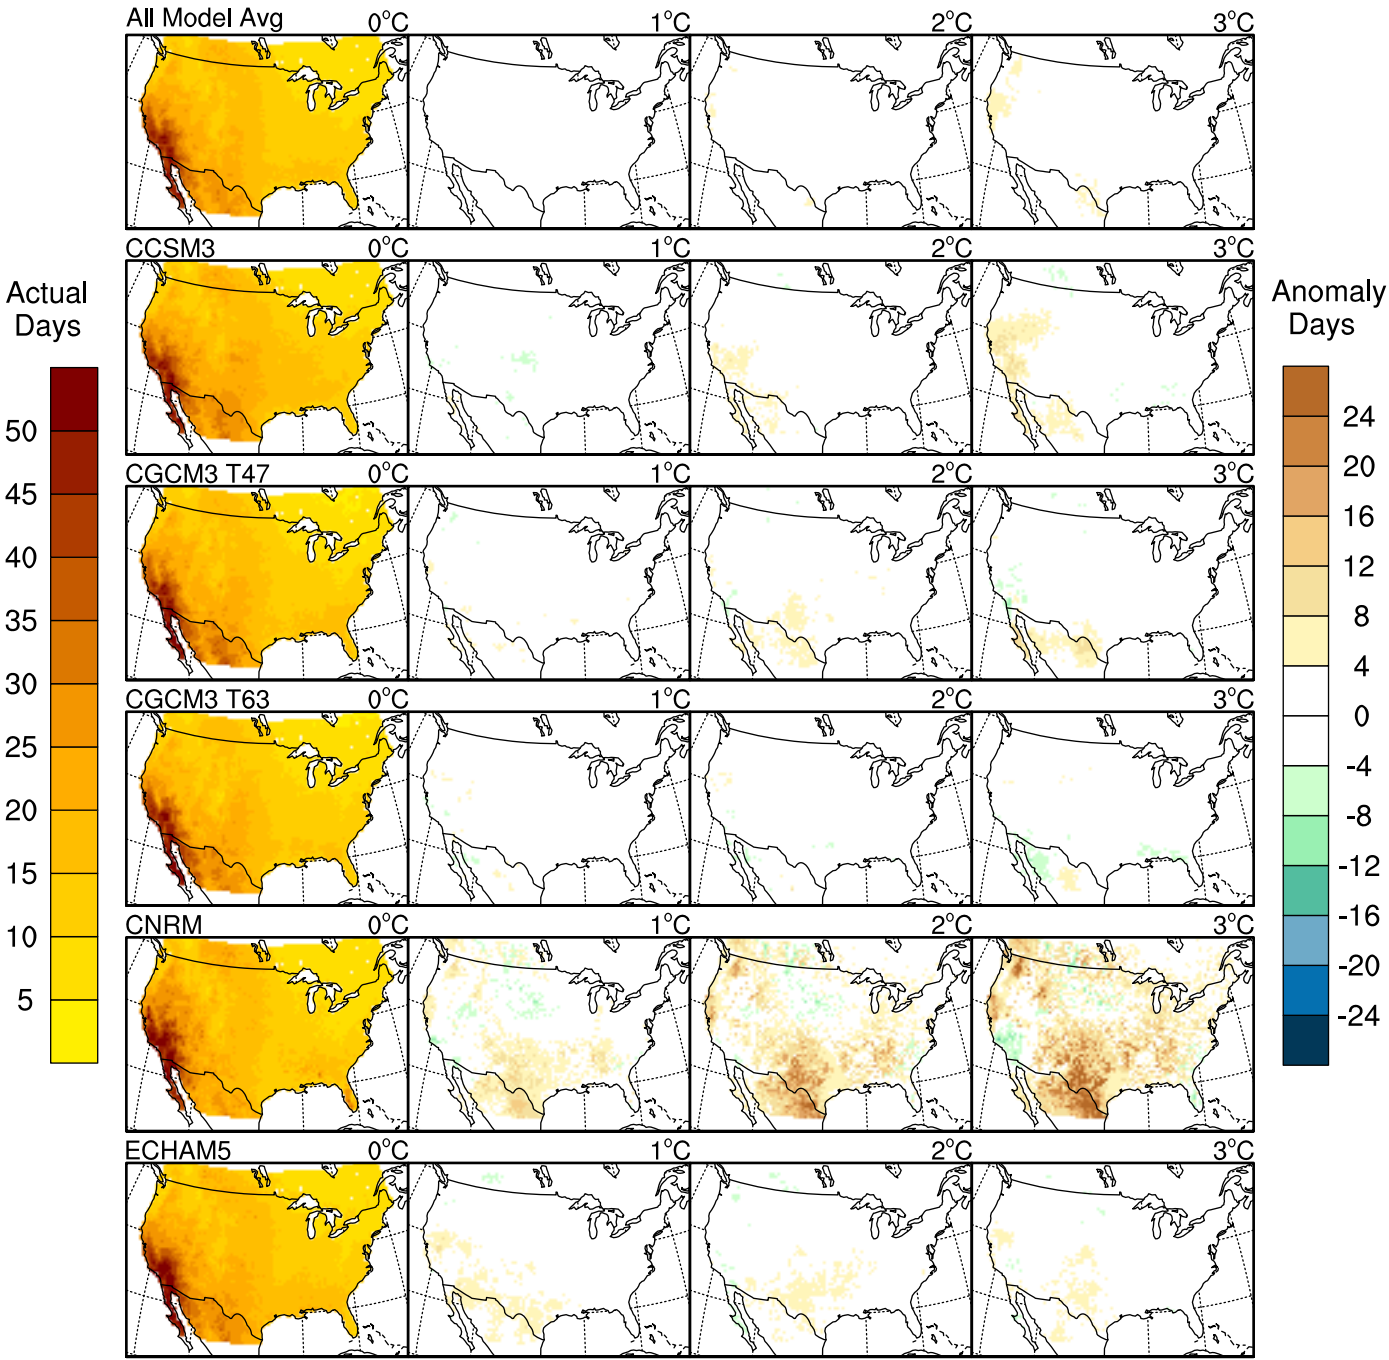
Longest Period of Dry Days < 0.01in (Summer)



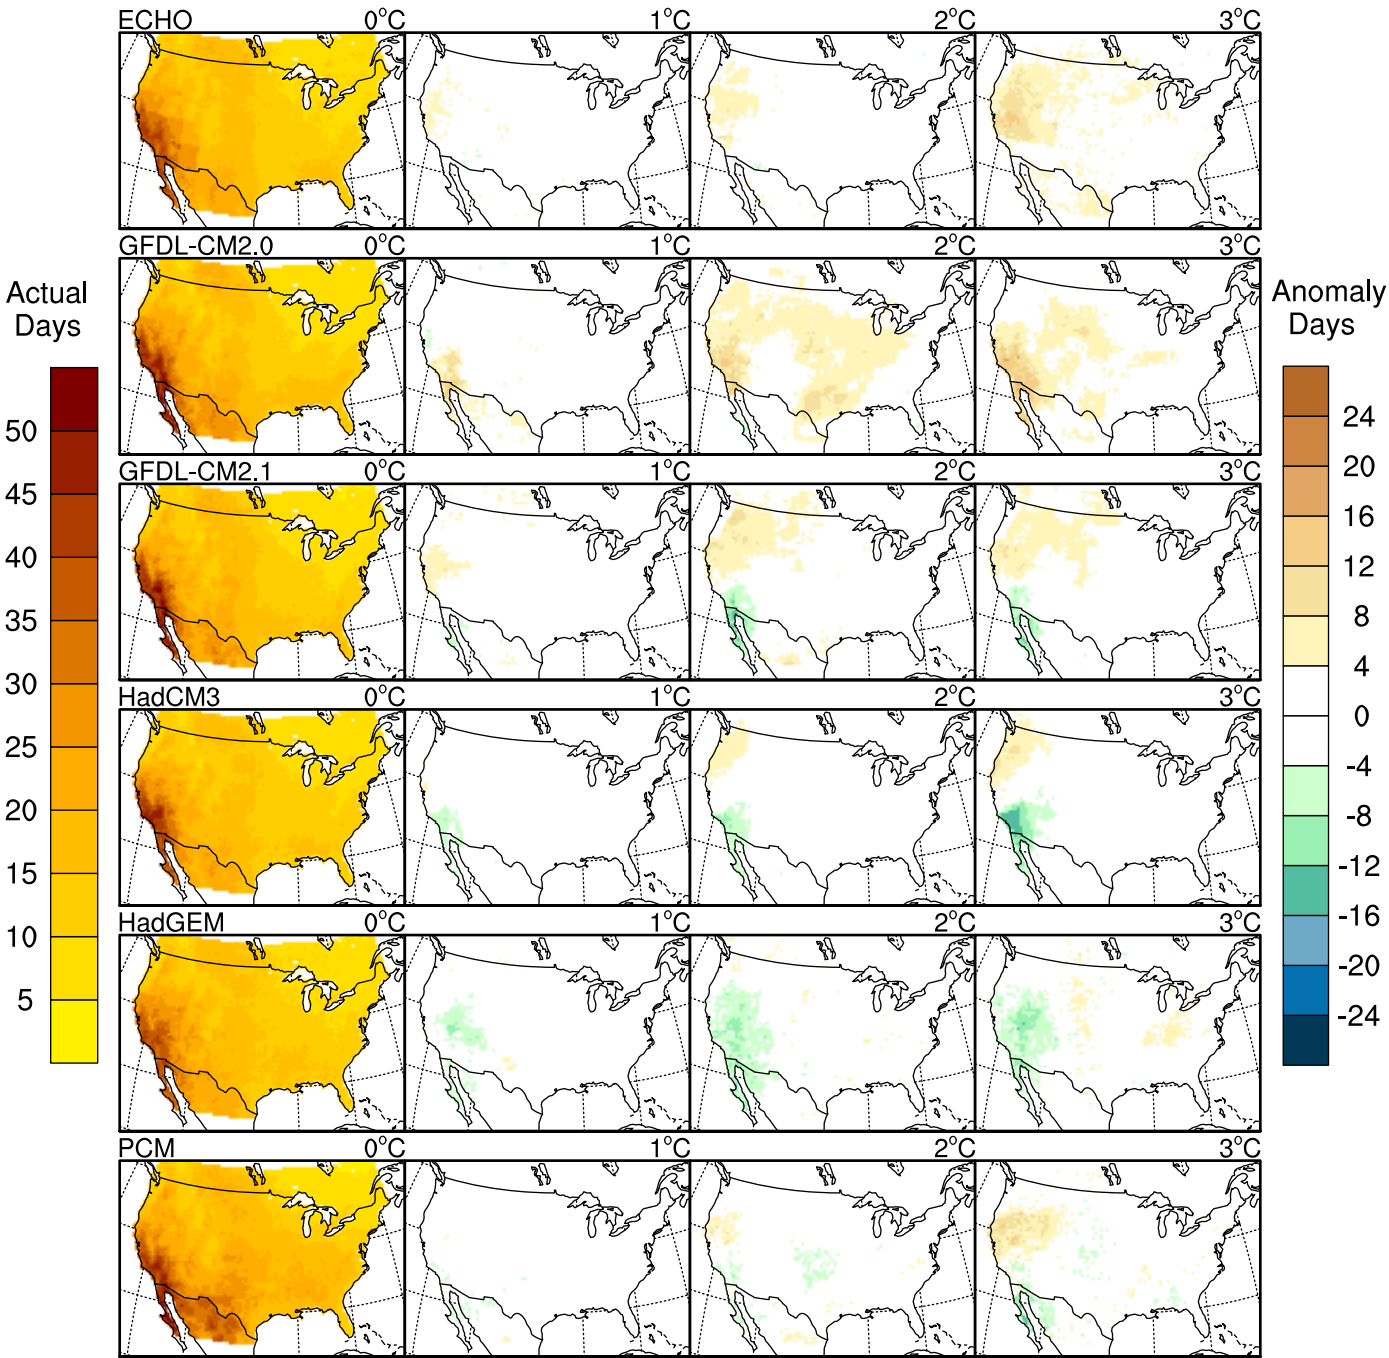
Longest Period of Dry Days < 0.01in (Summer)



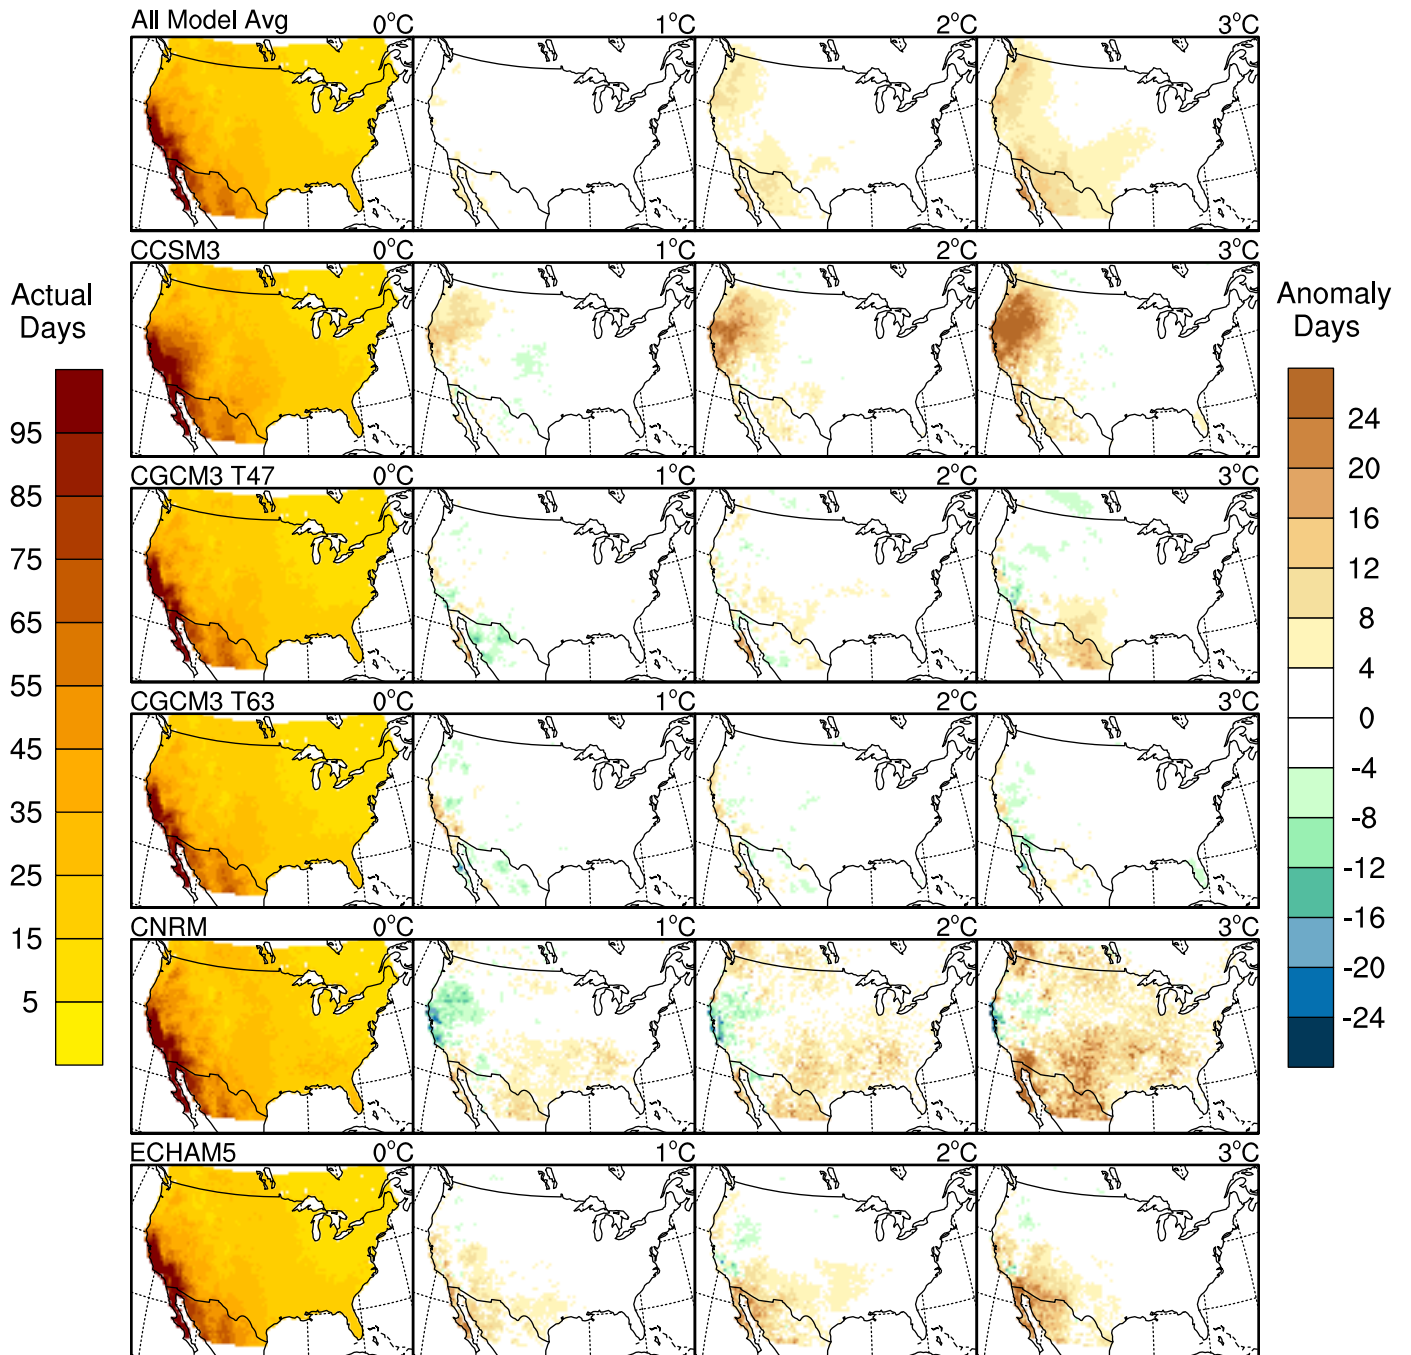
Longest Period of Dry Days < 0.01in (Fall)



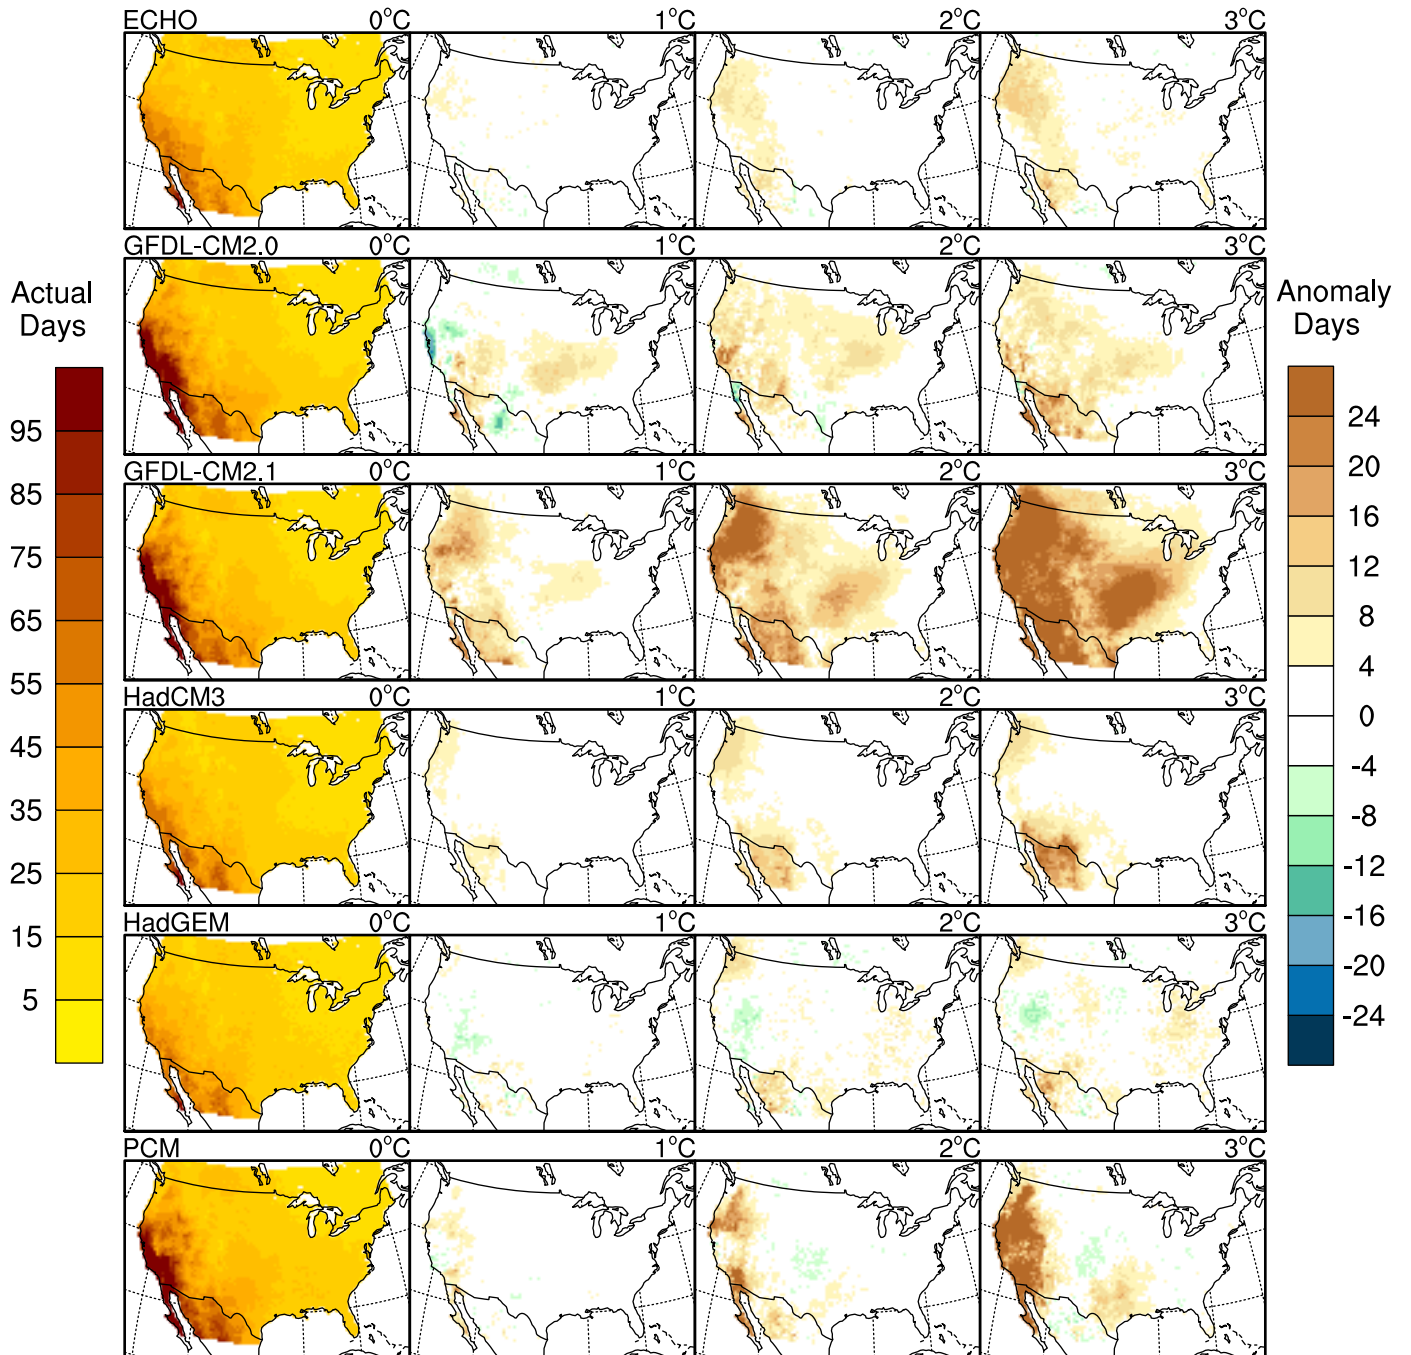
Longest Period of Dry Days < 0.01in (Fall)



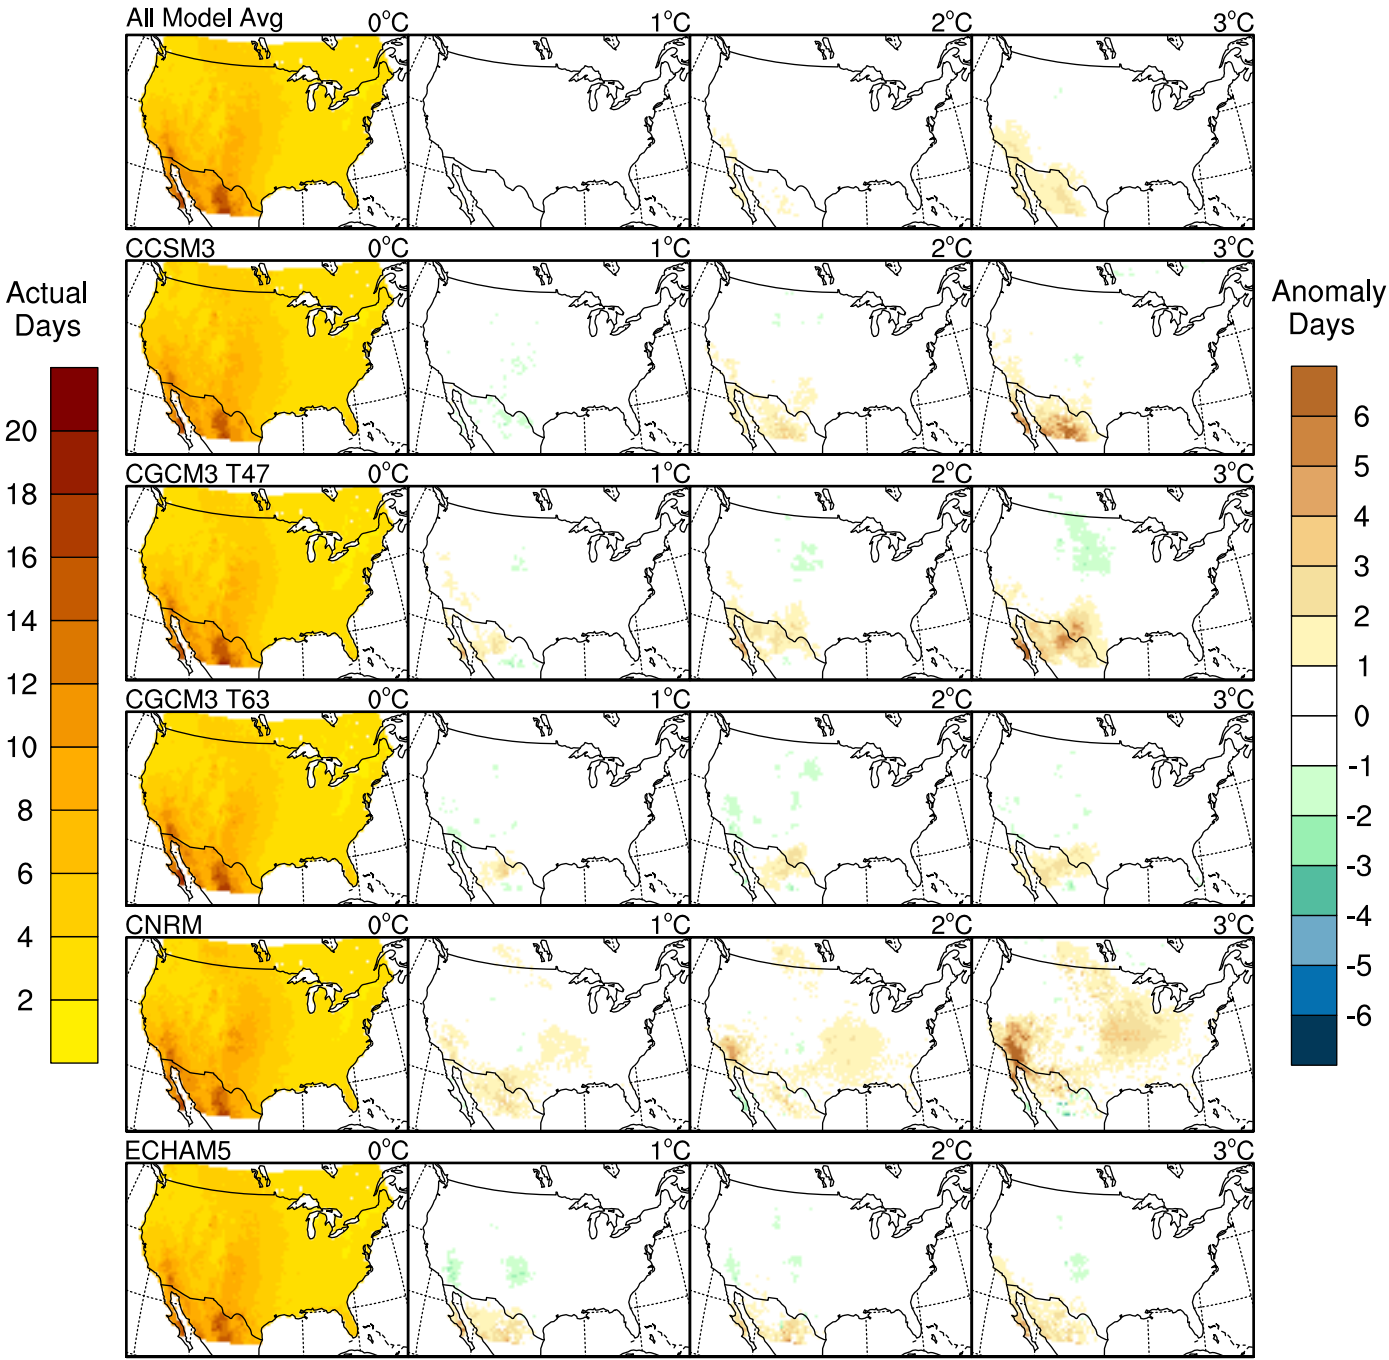
Longest Period of Dry Days < 0.01in (Annual)



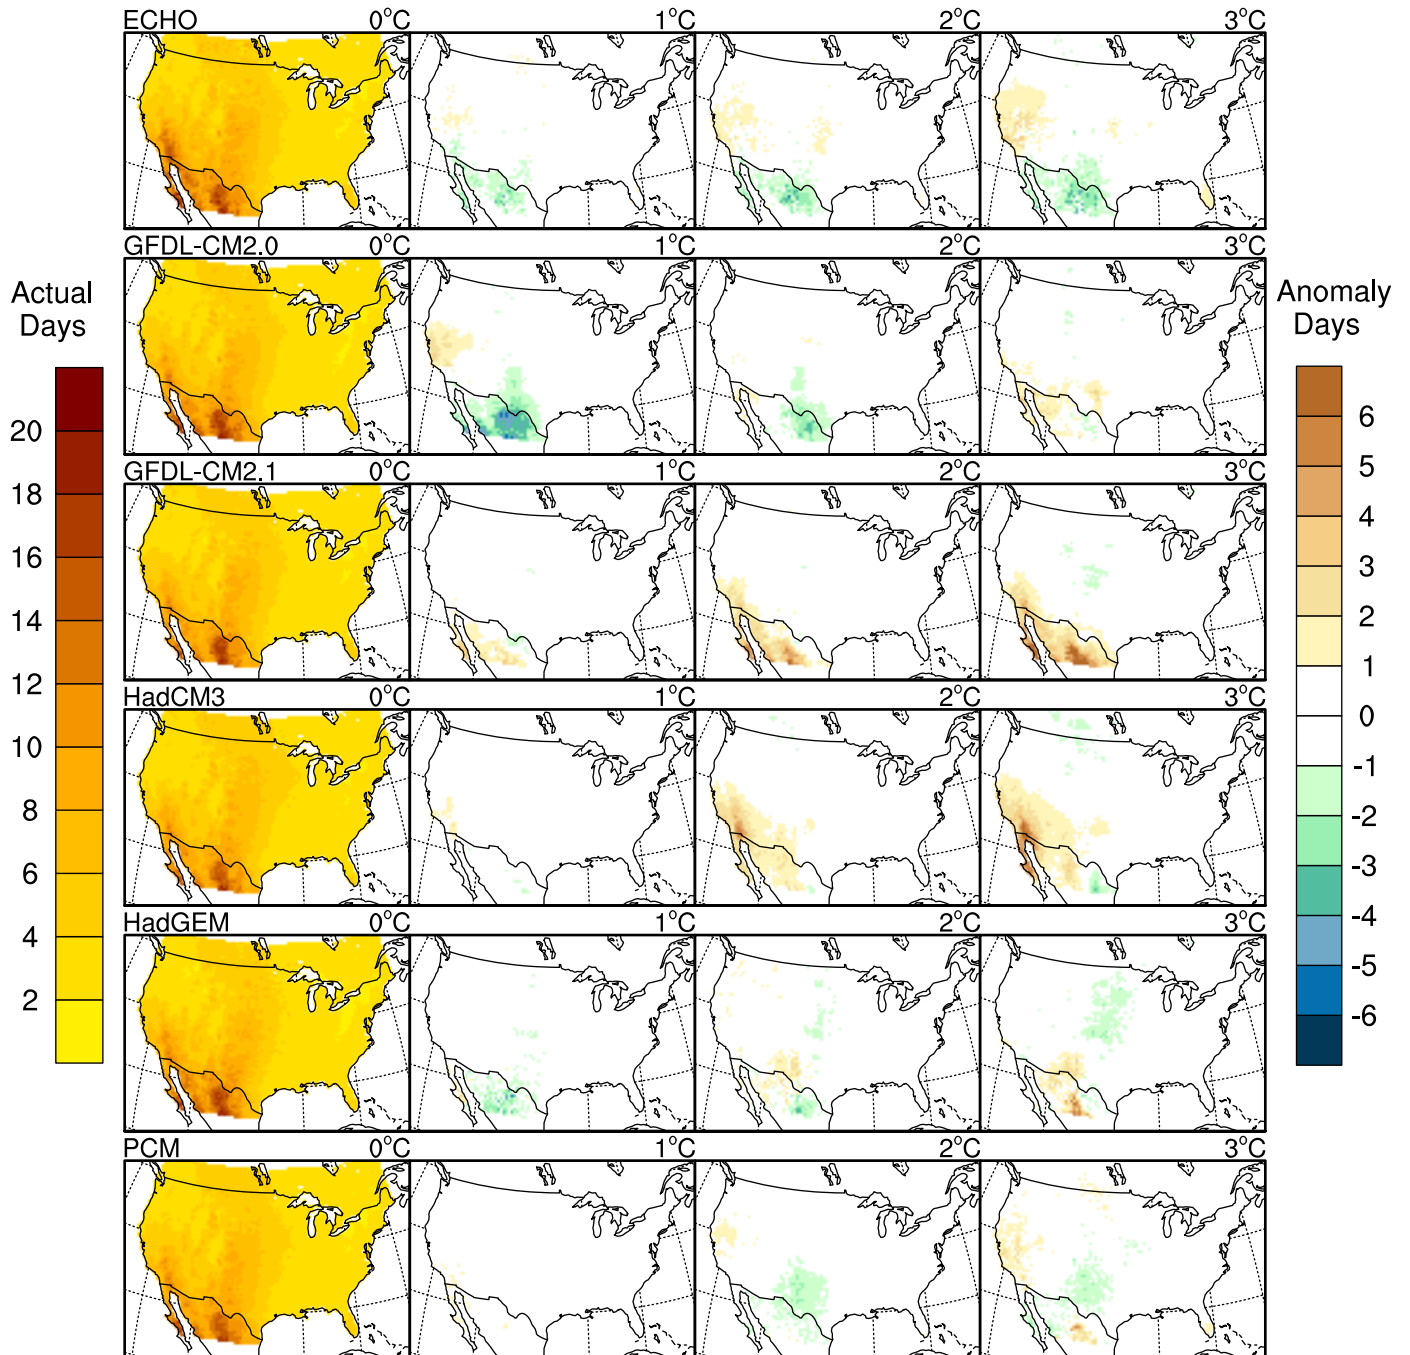
Longest Period of Dry Days < 0.01in (Annual)



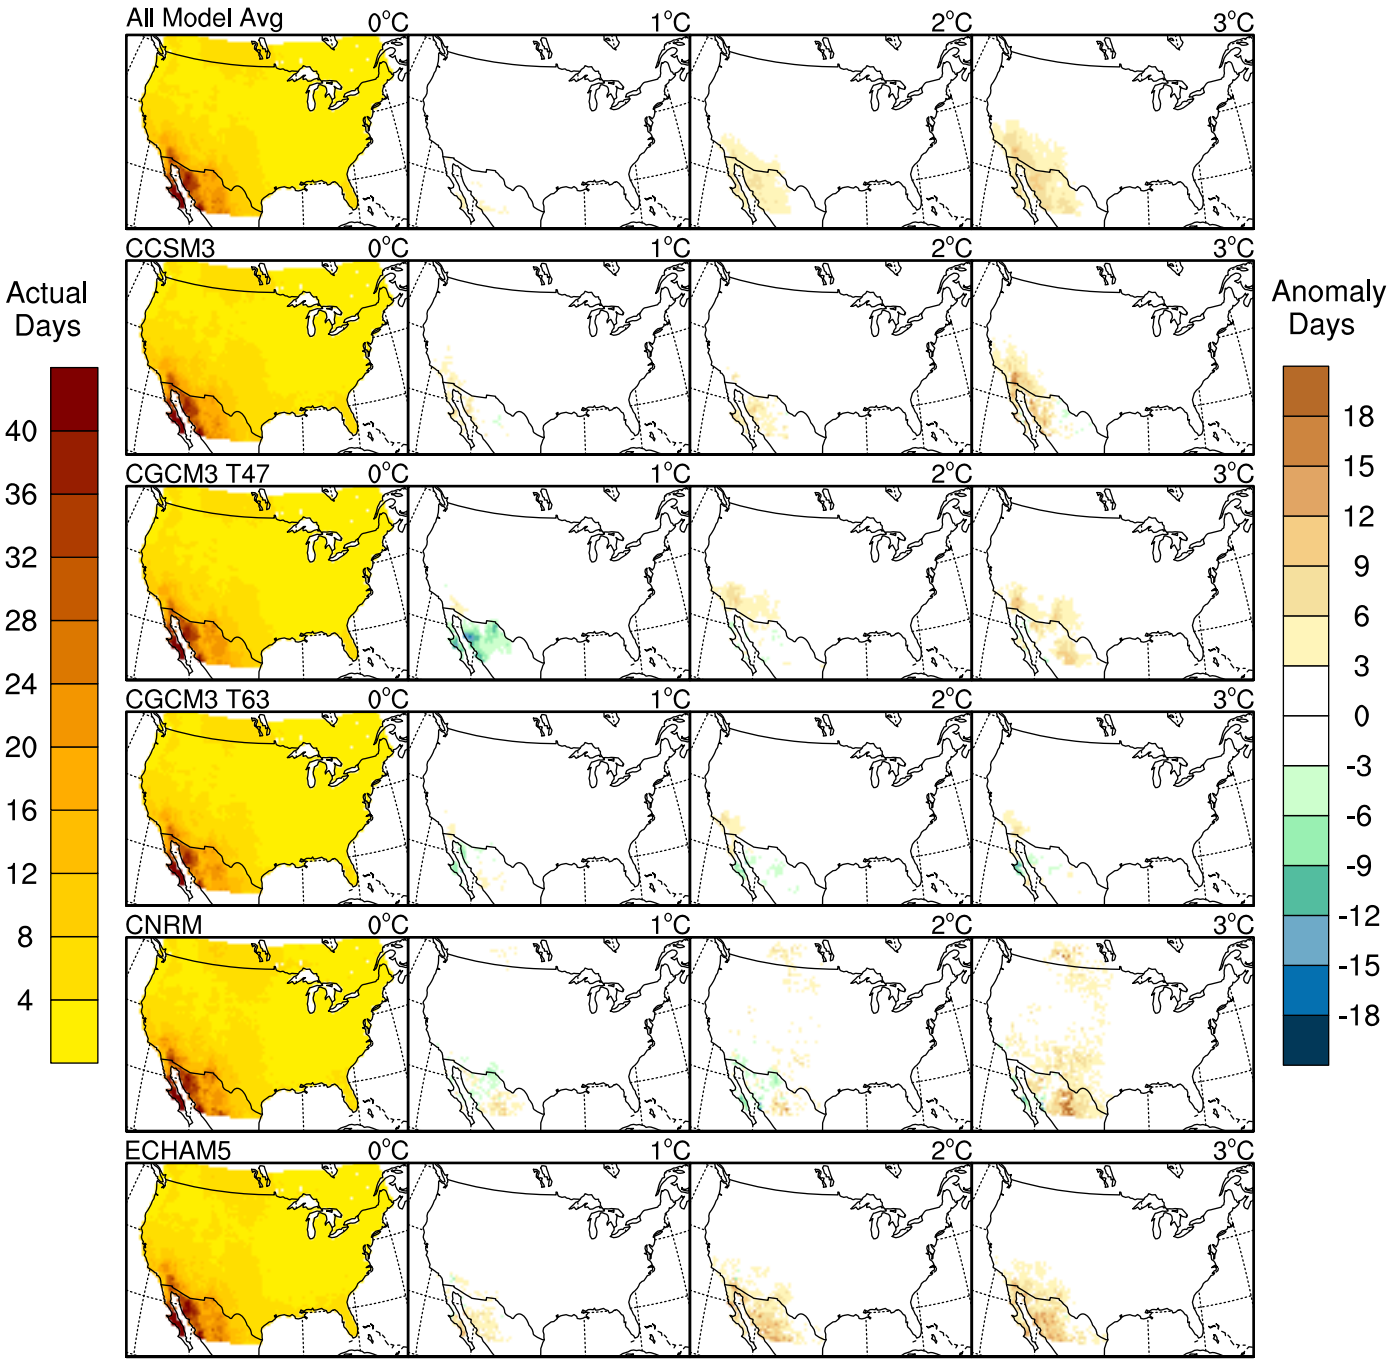
Avg Length of Dry Periods < 0.01in/day (Winter)



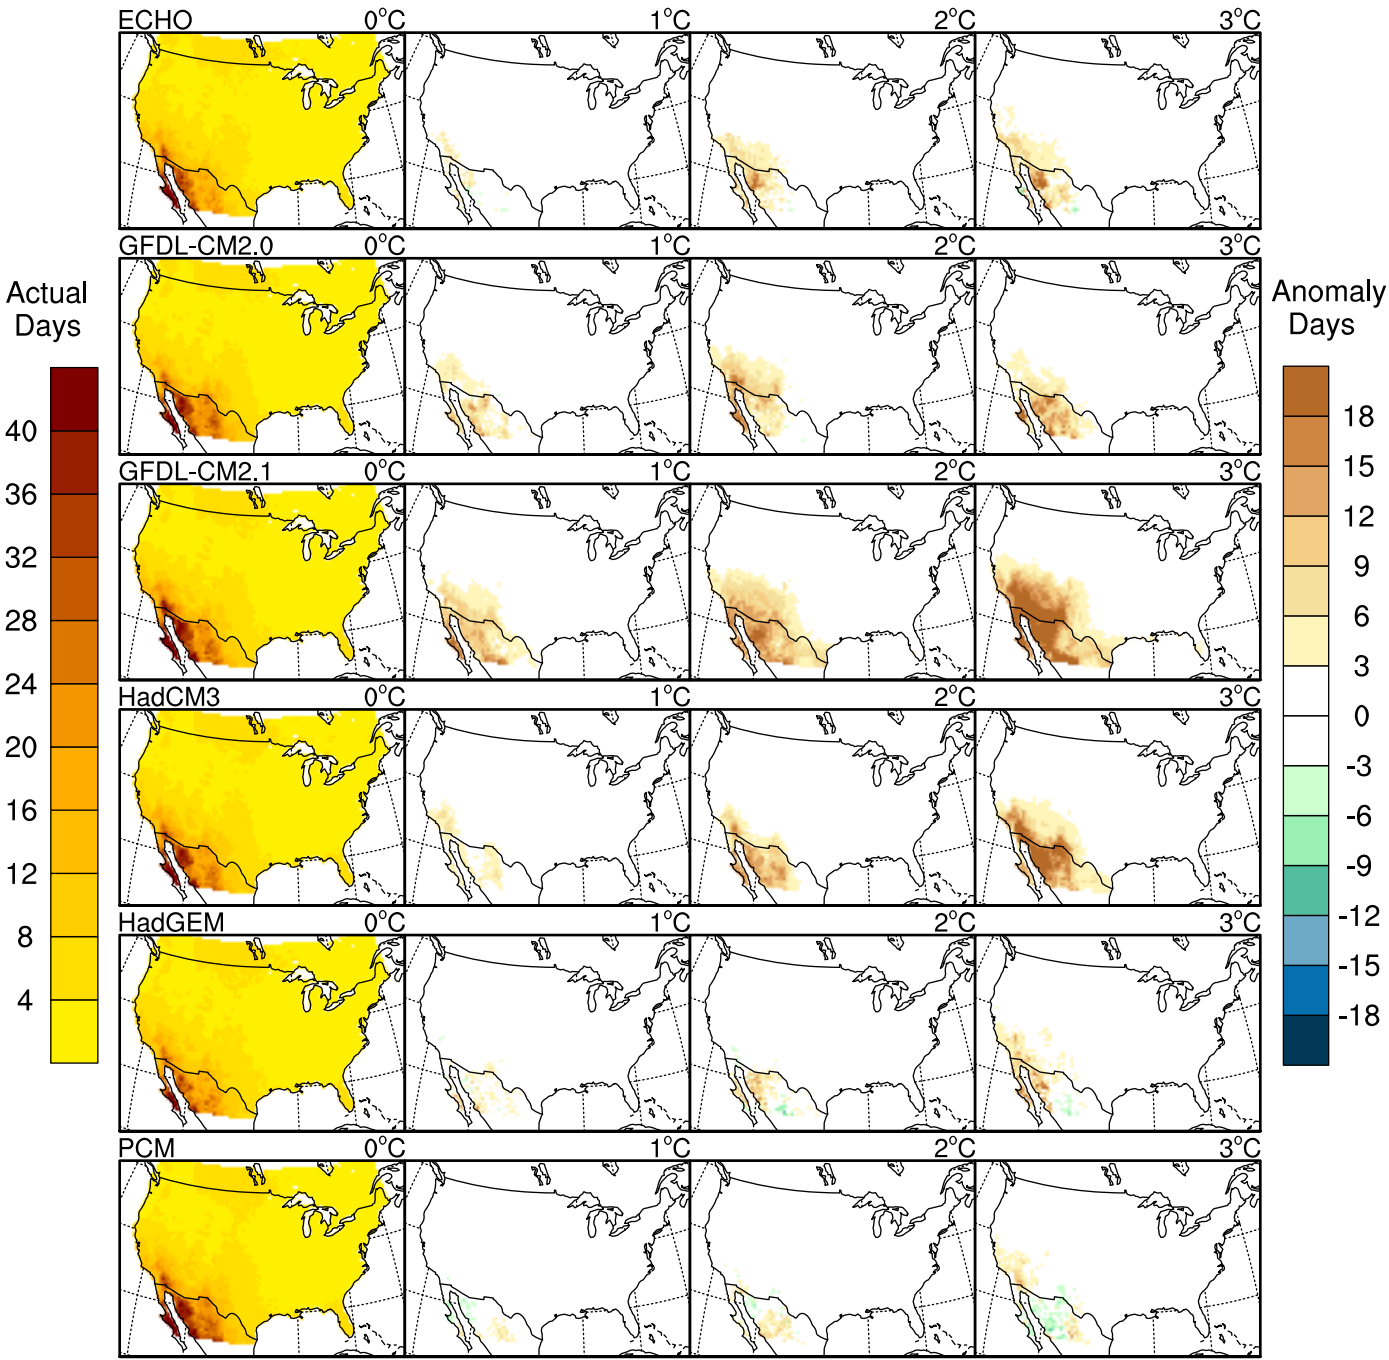
Avg Length of Dry Periods < 0.01in/day (Winter)



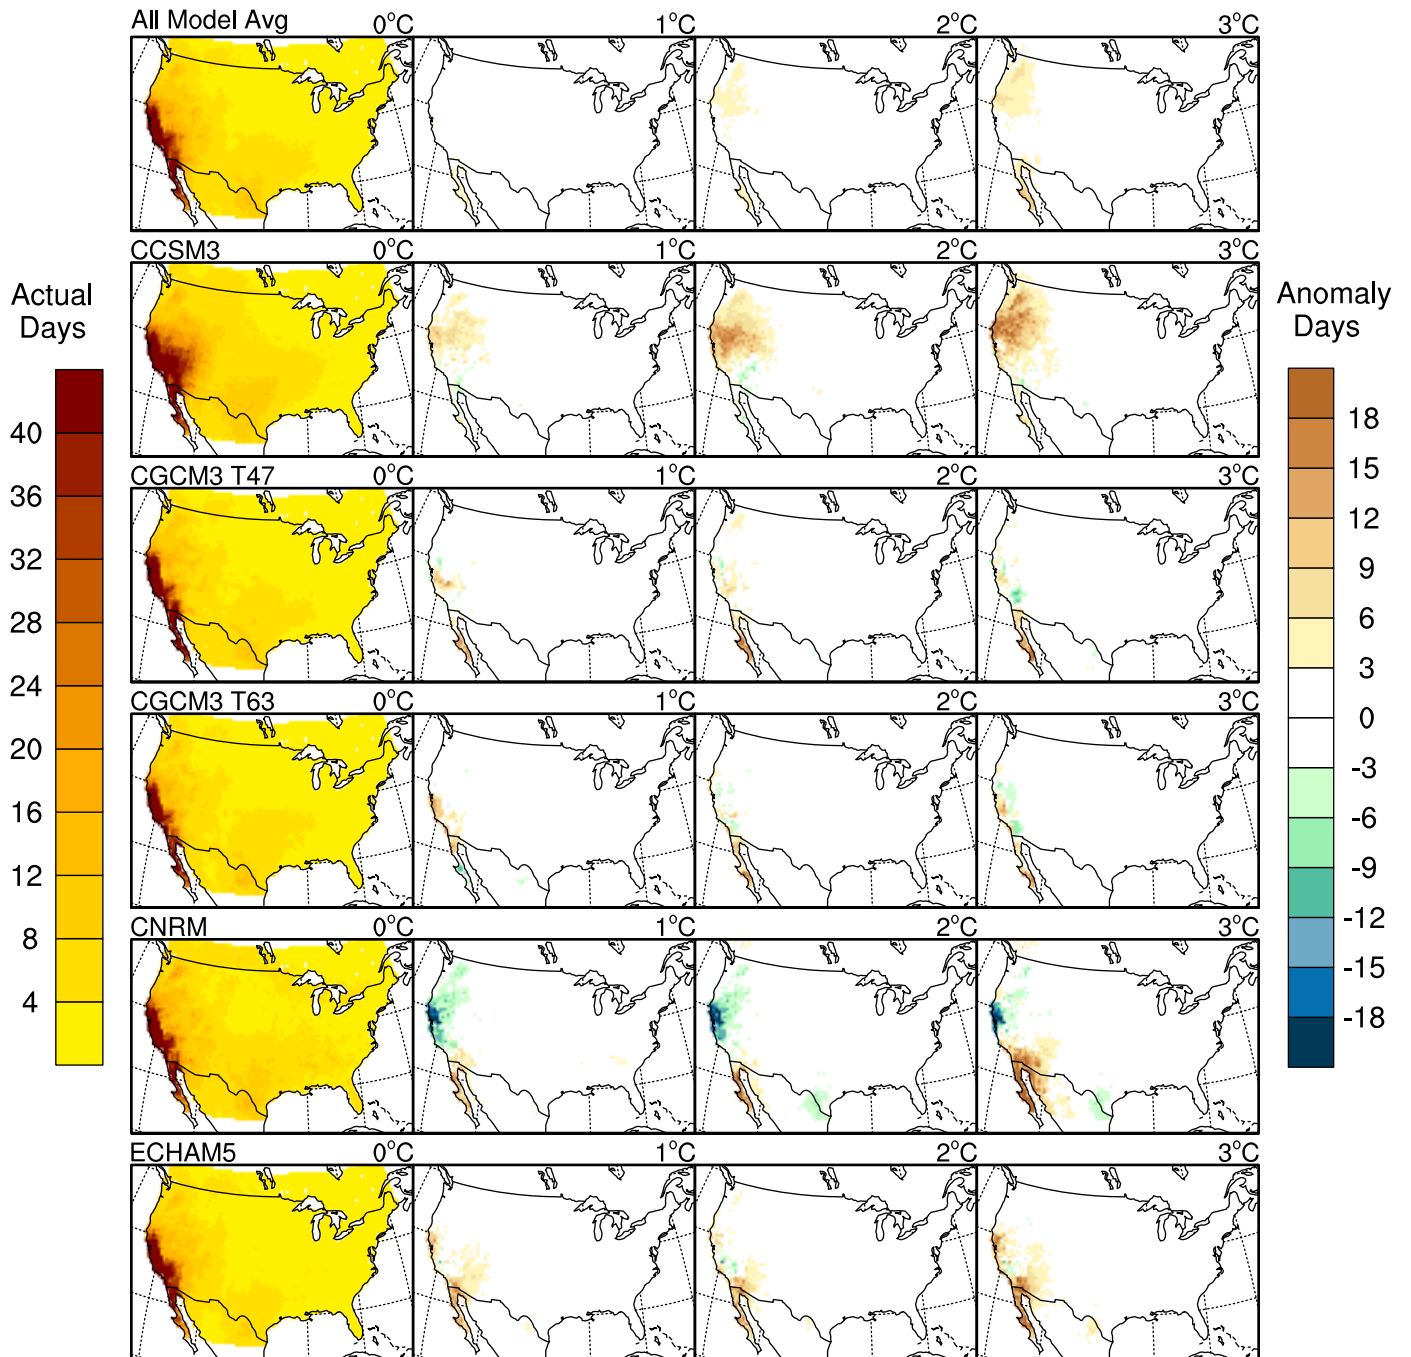
Avg Length of Dry Periods < 0.01in/day (Spring)



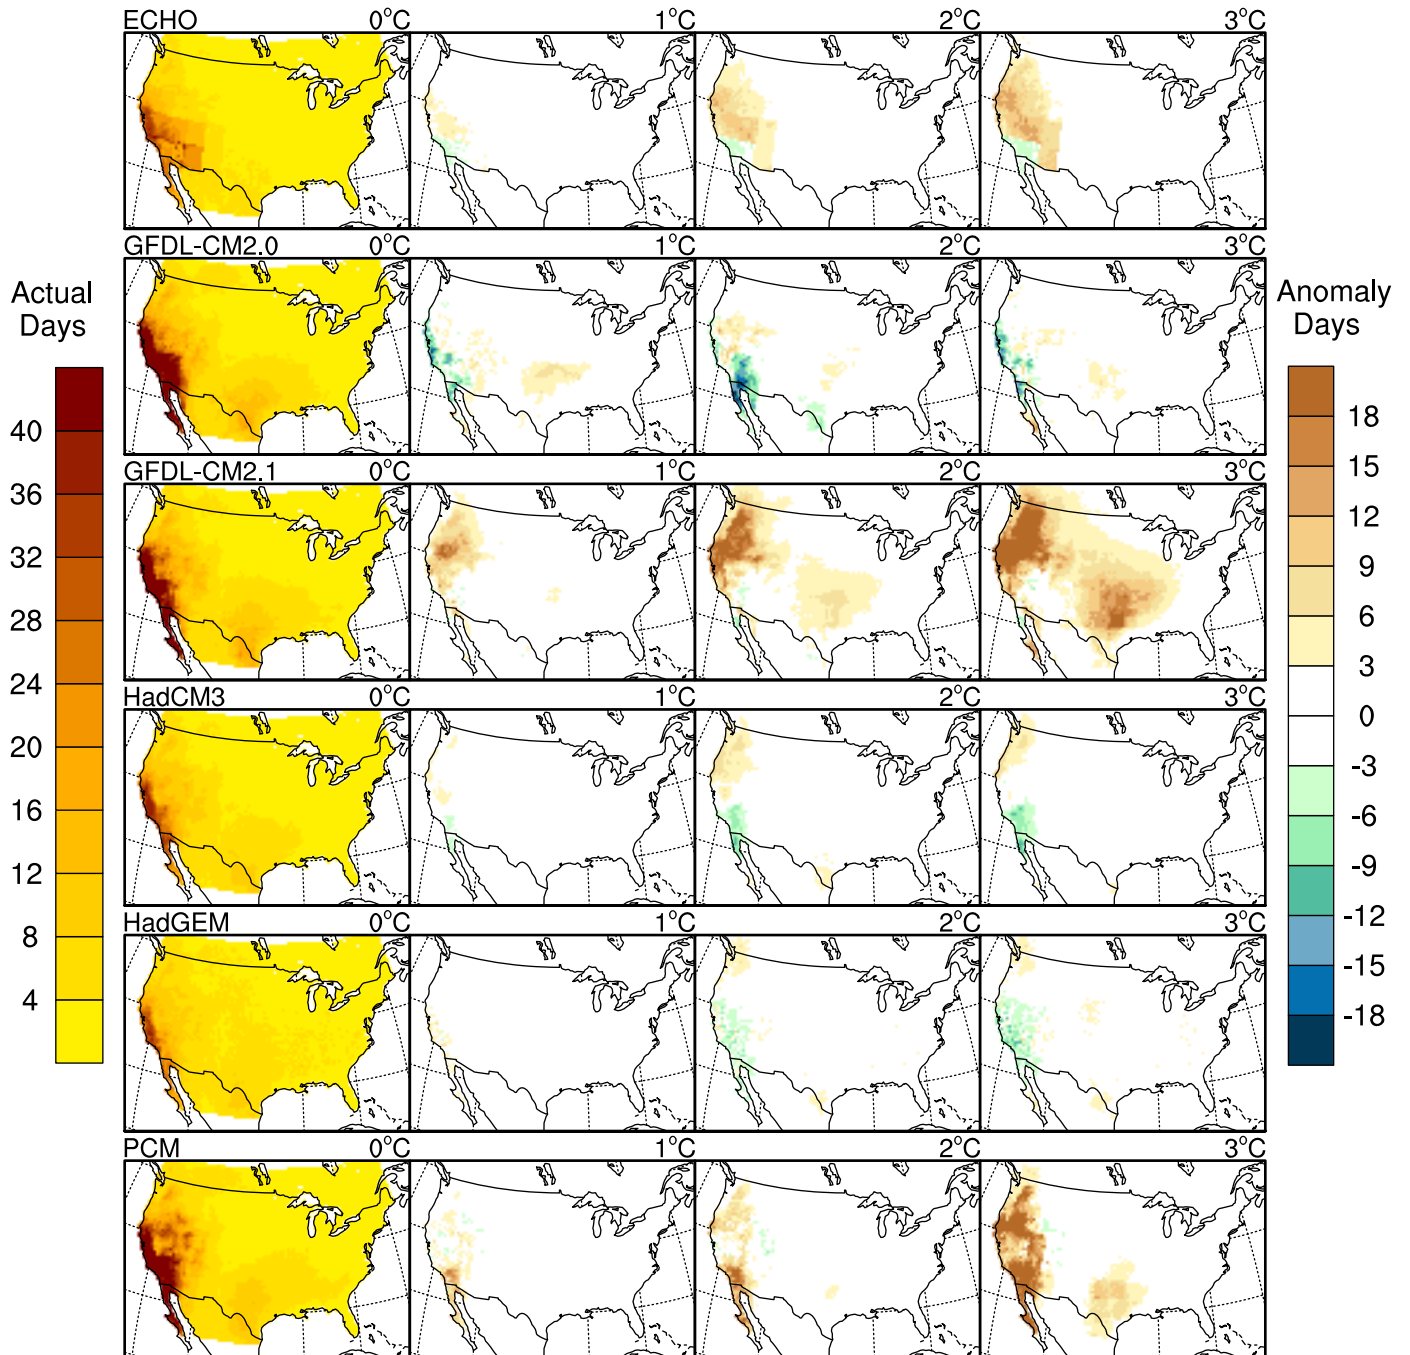
Avg Length of Dry Periods < 0.01in/day (Spring)



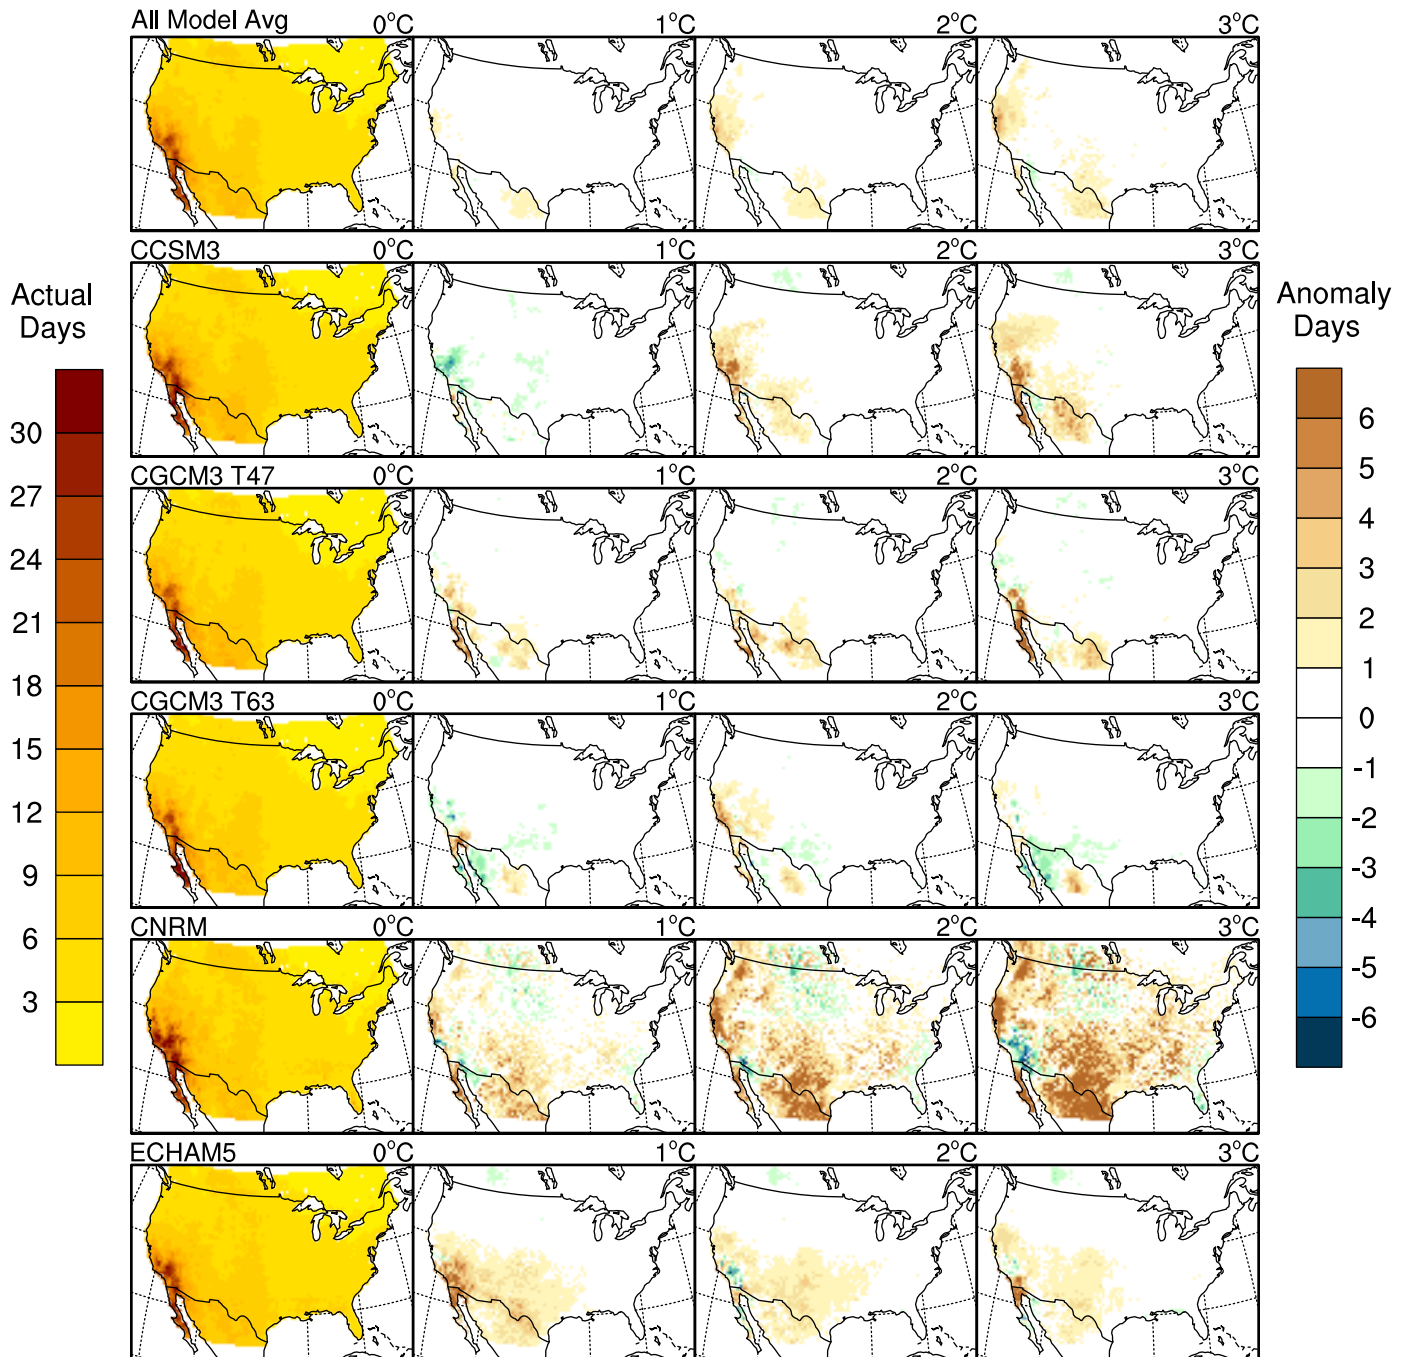
Avg Length of Dry Periods < 0.01in/day (Summer)



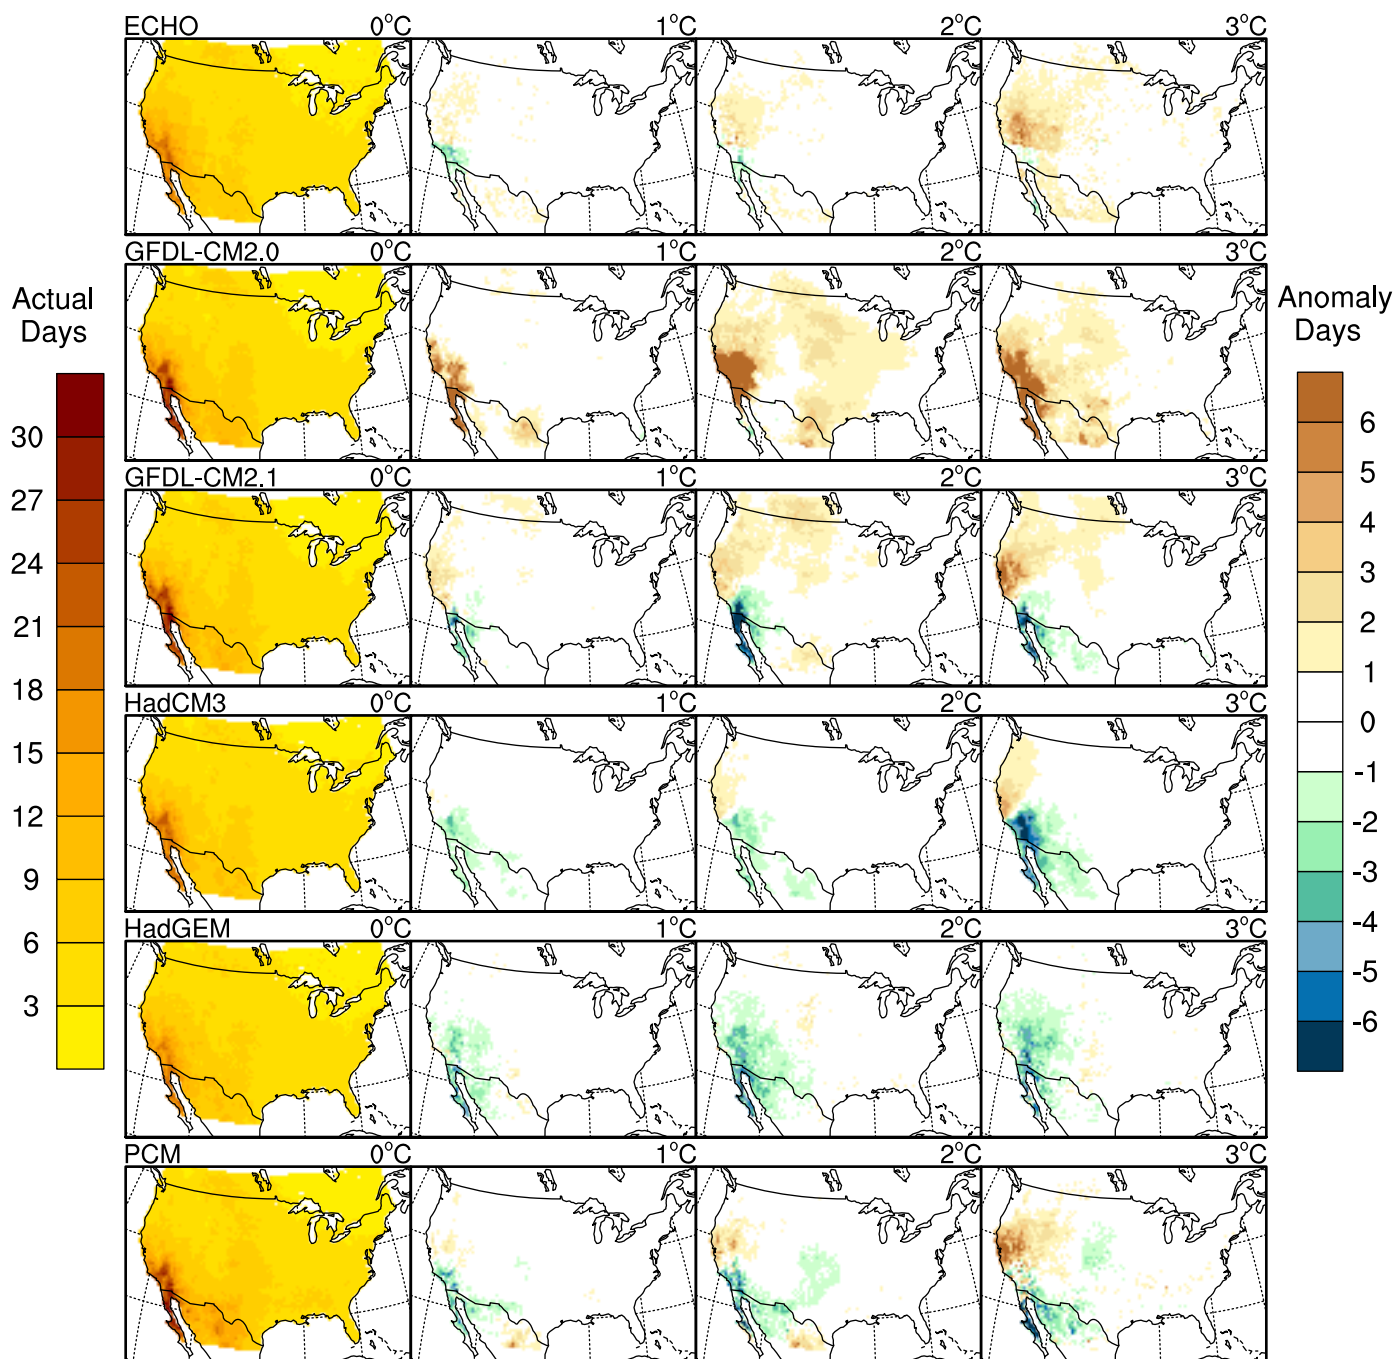
Avg Length of Dry Periods < 0.01in/day (Summer)



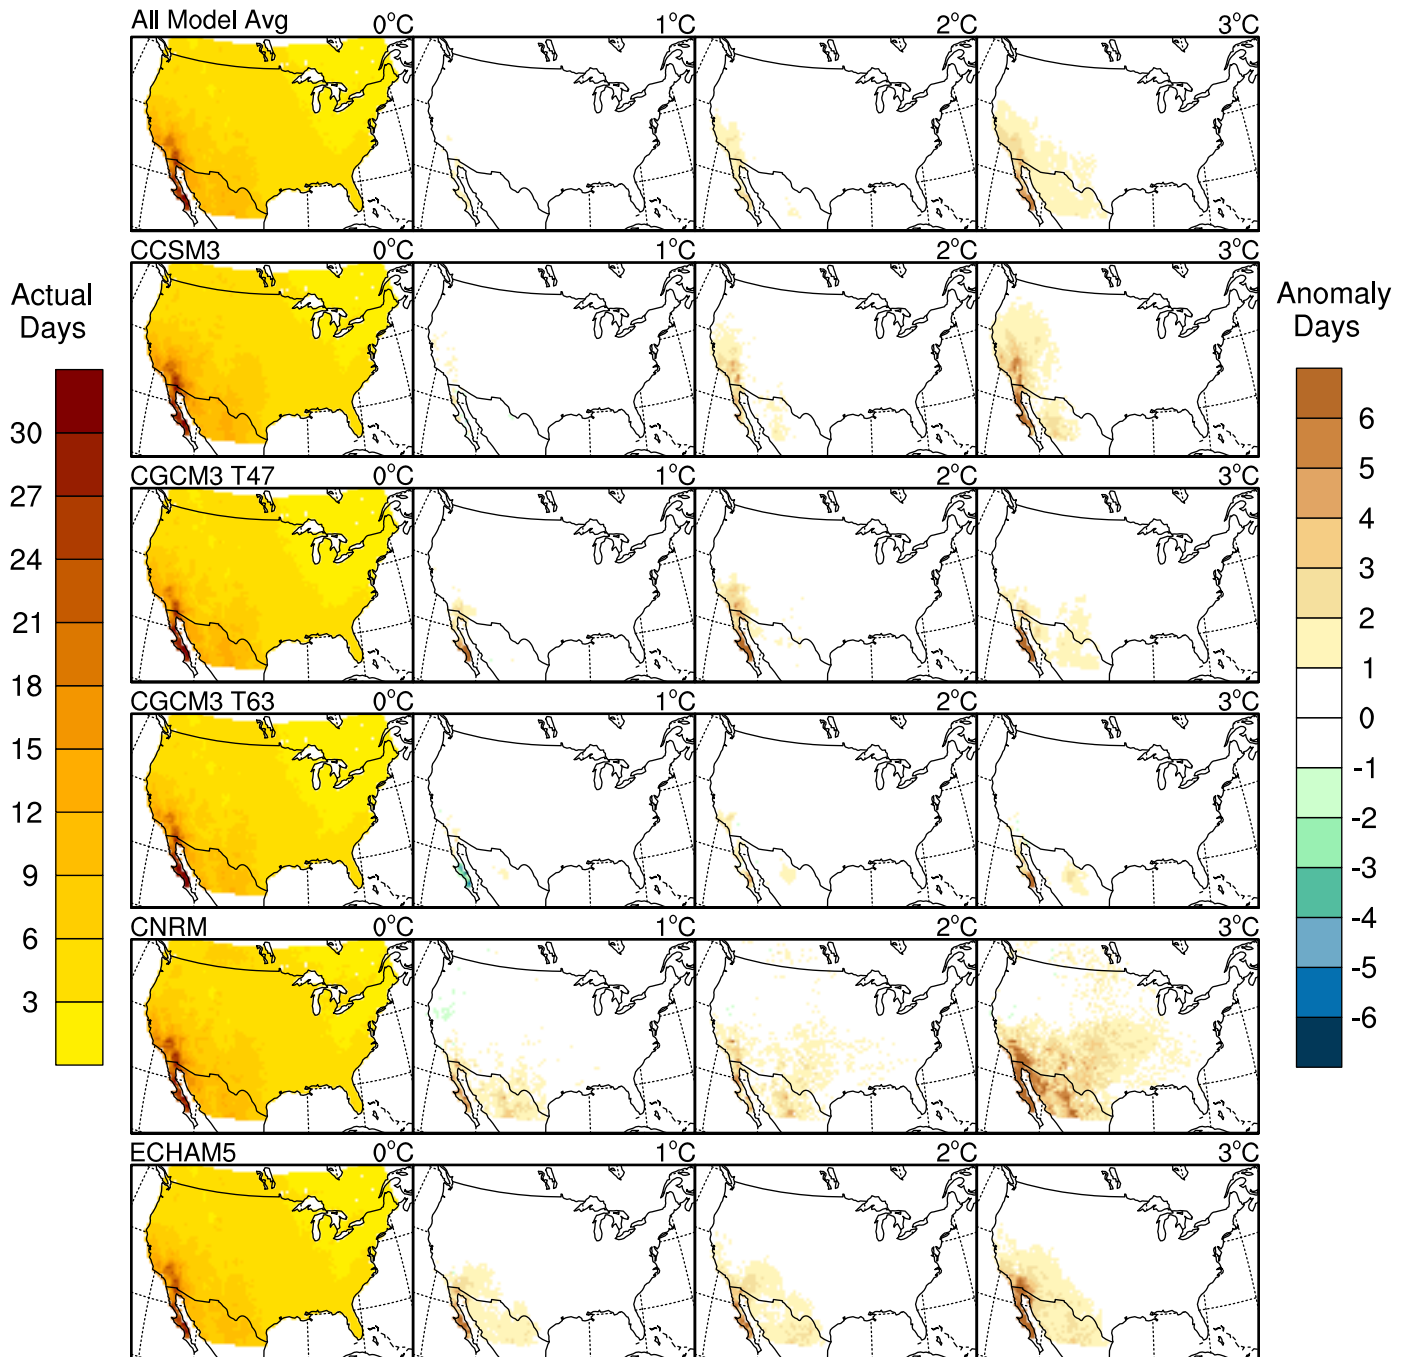
Avg Length of Dry Periods < 0.01in/day (Fall)



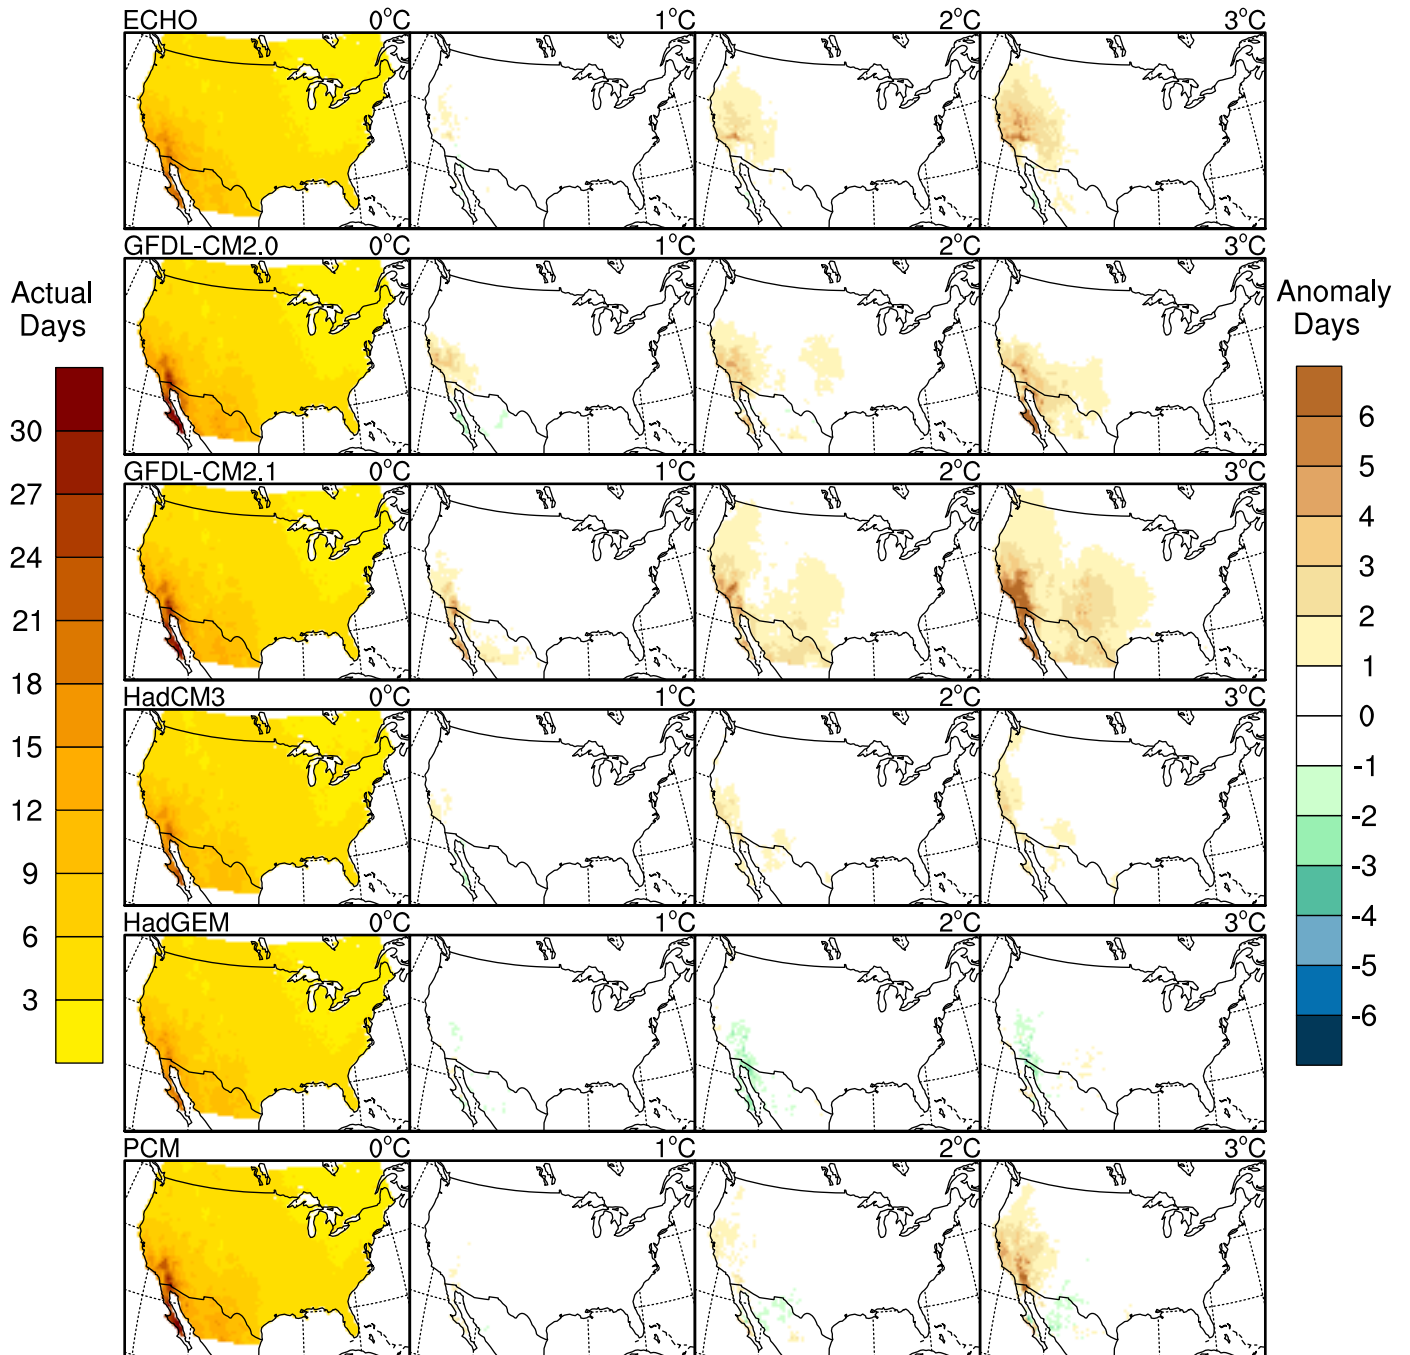
Avg Length of Dry Periods < 0.01in/day (Fall)



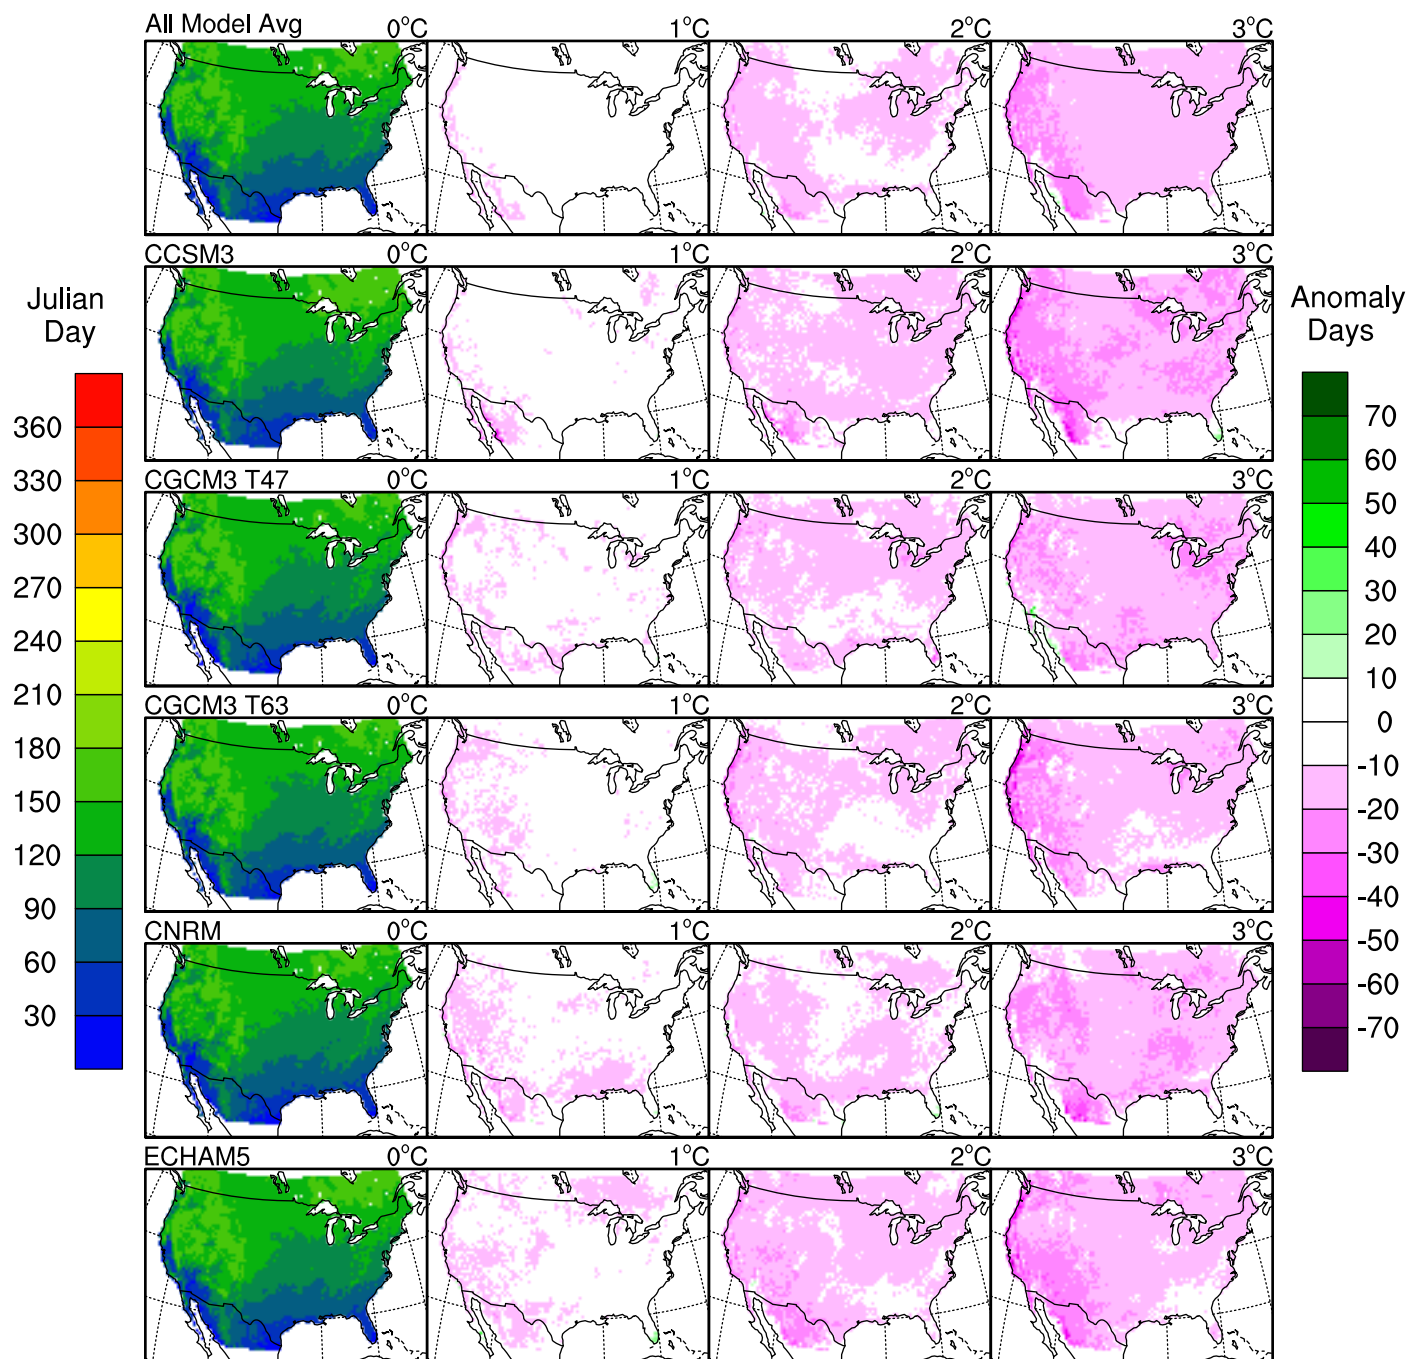
Avg Length of Dry Periods < 0.01in/day (Annual)



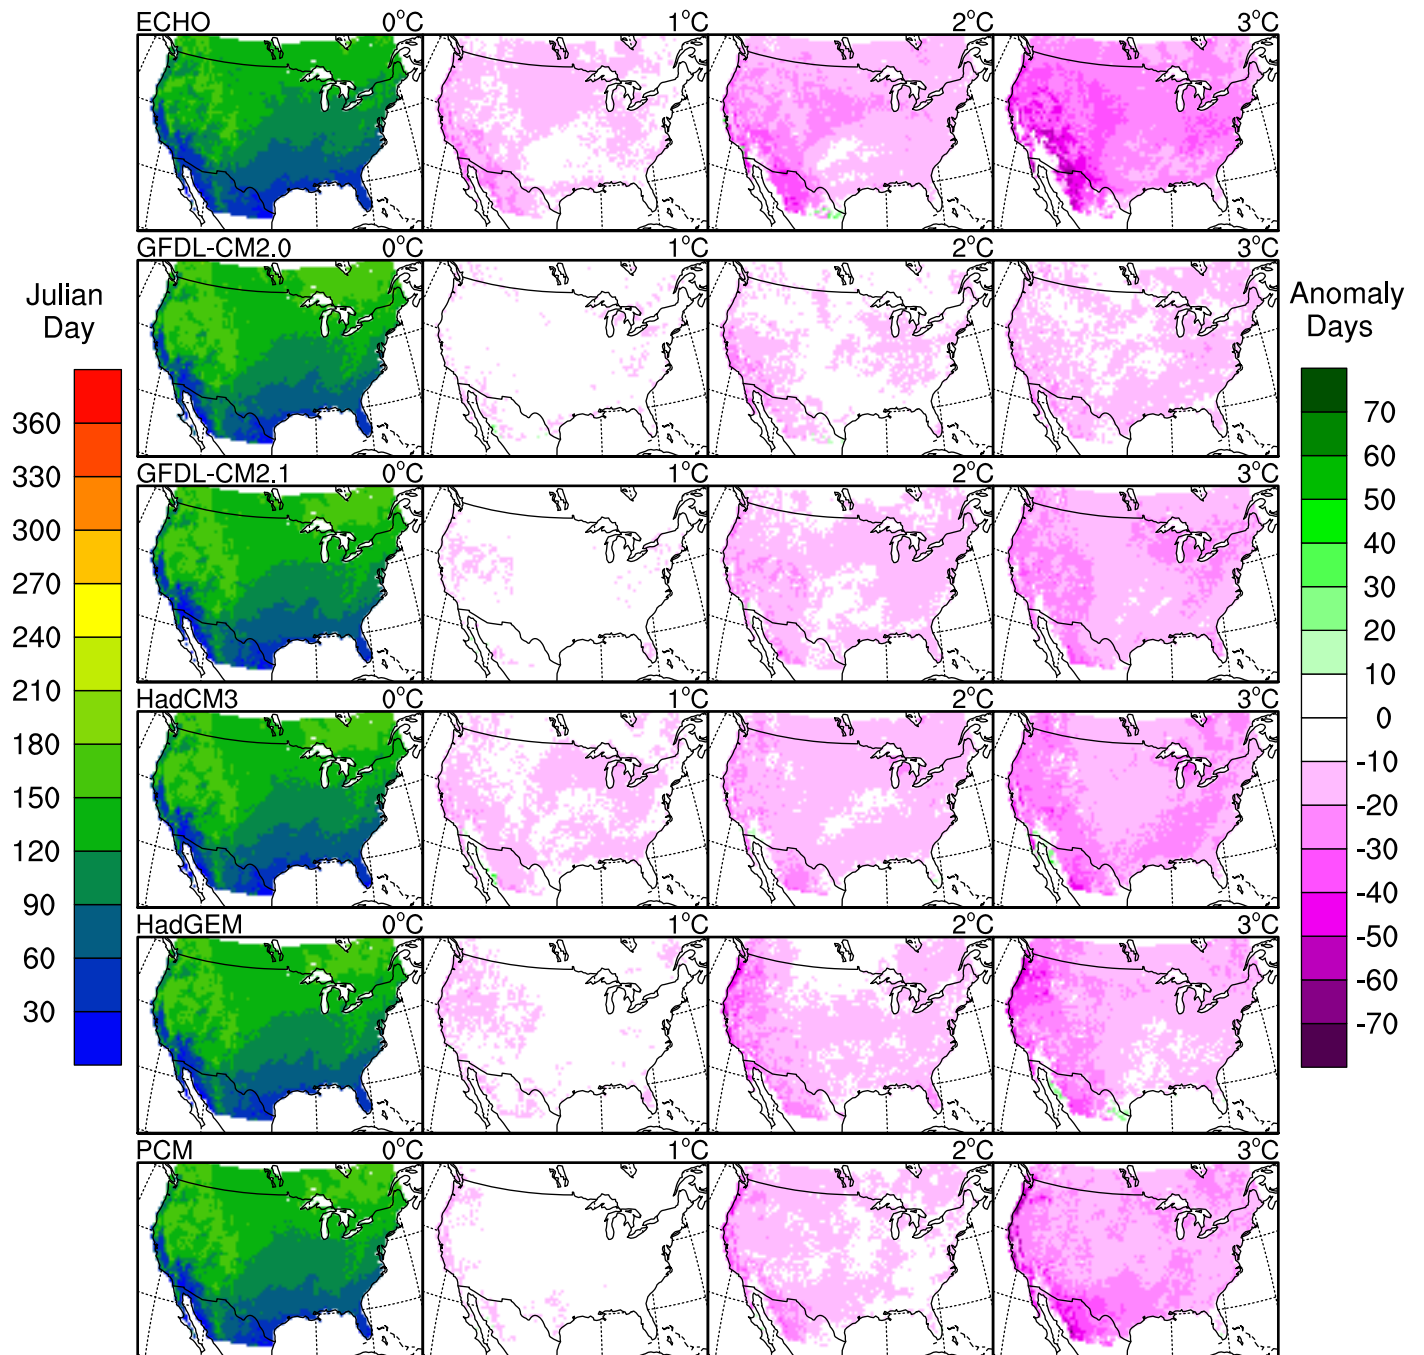
Avg Length of Dry Periods < 0.01in/day (Annual)



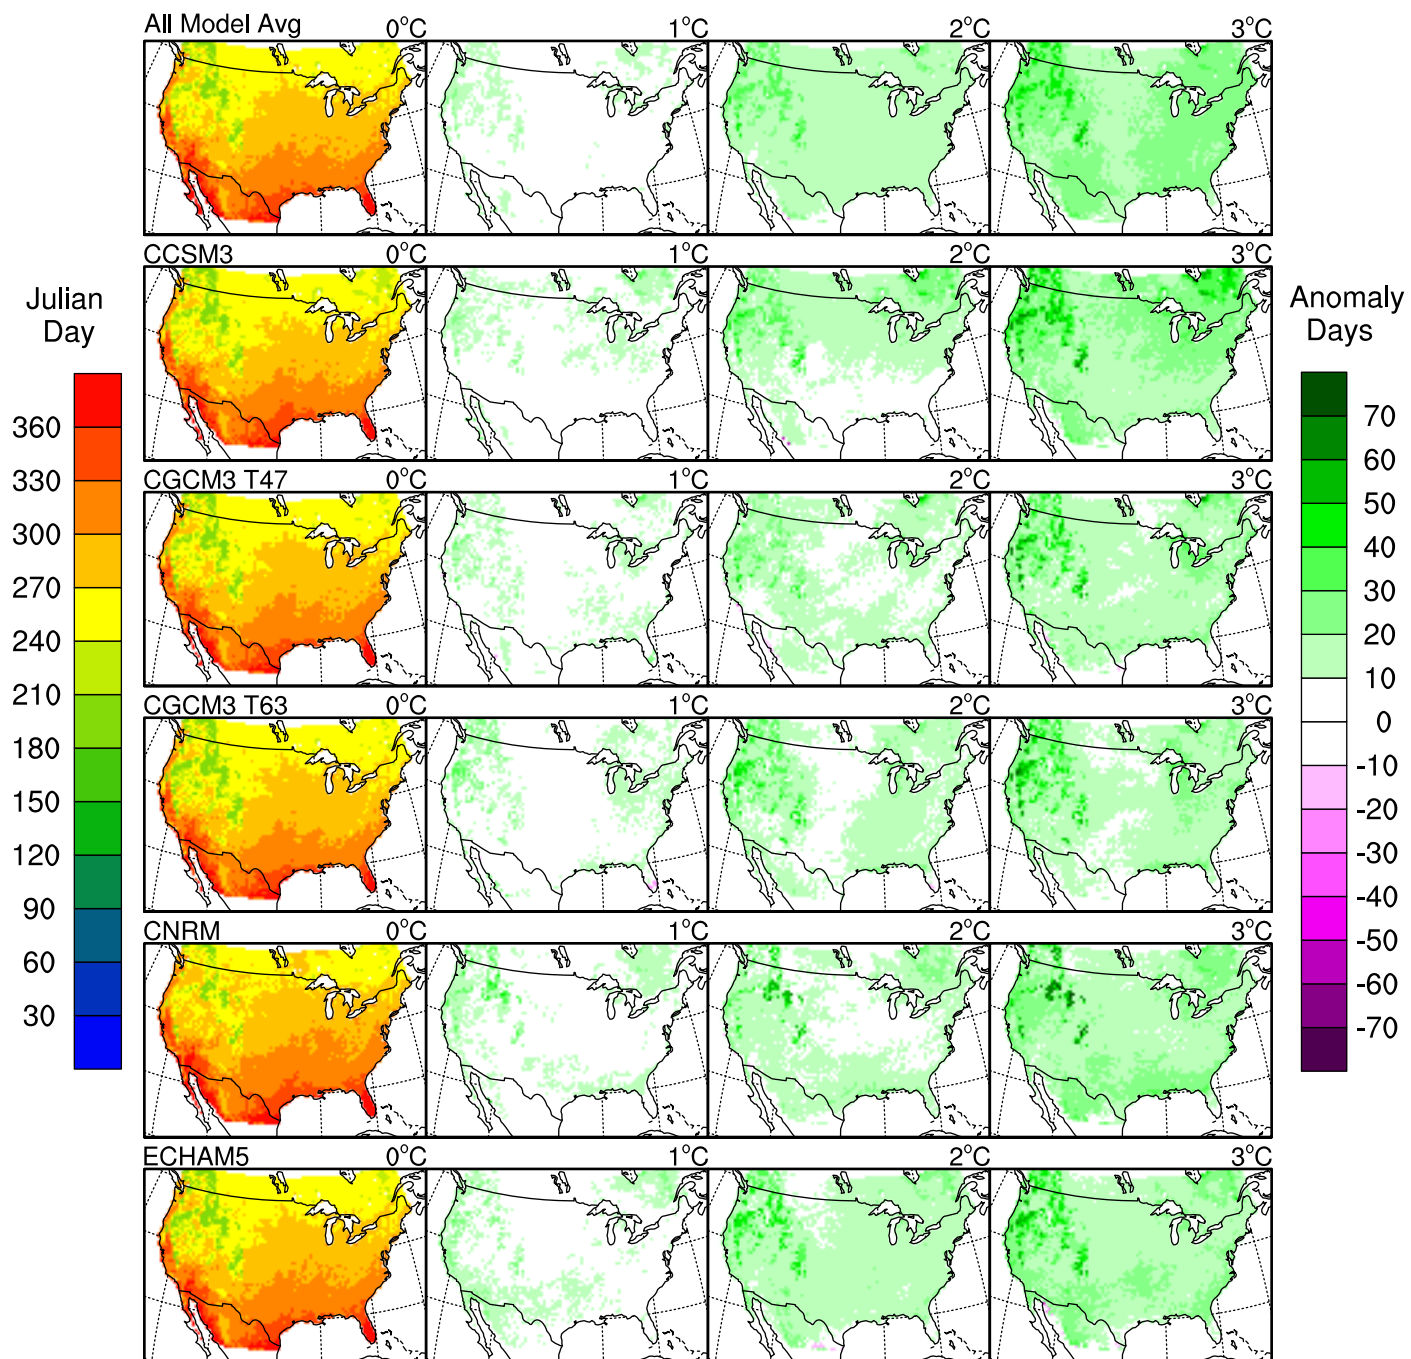
Beginning of Growing Season (Julian Day)



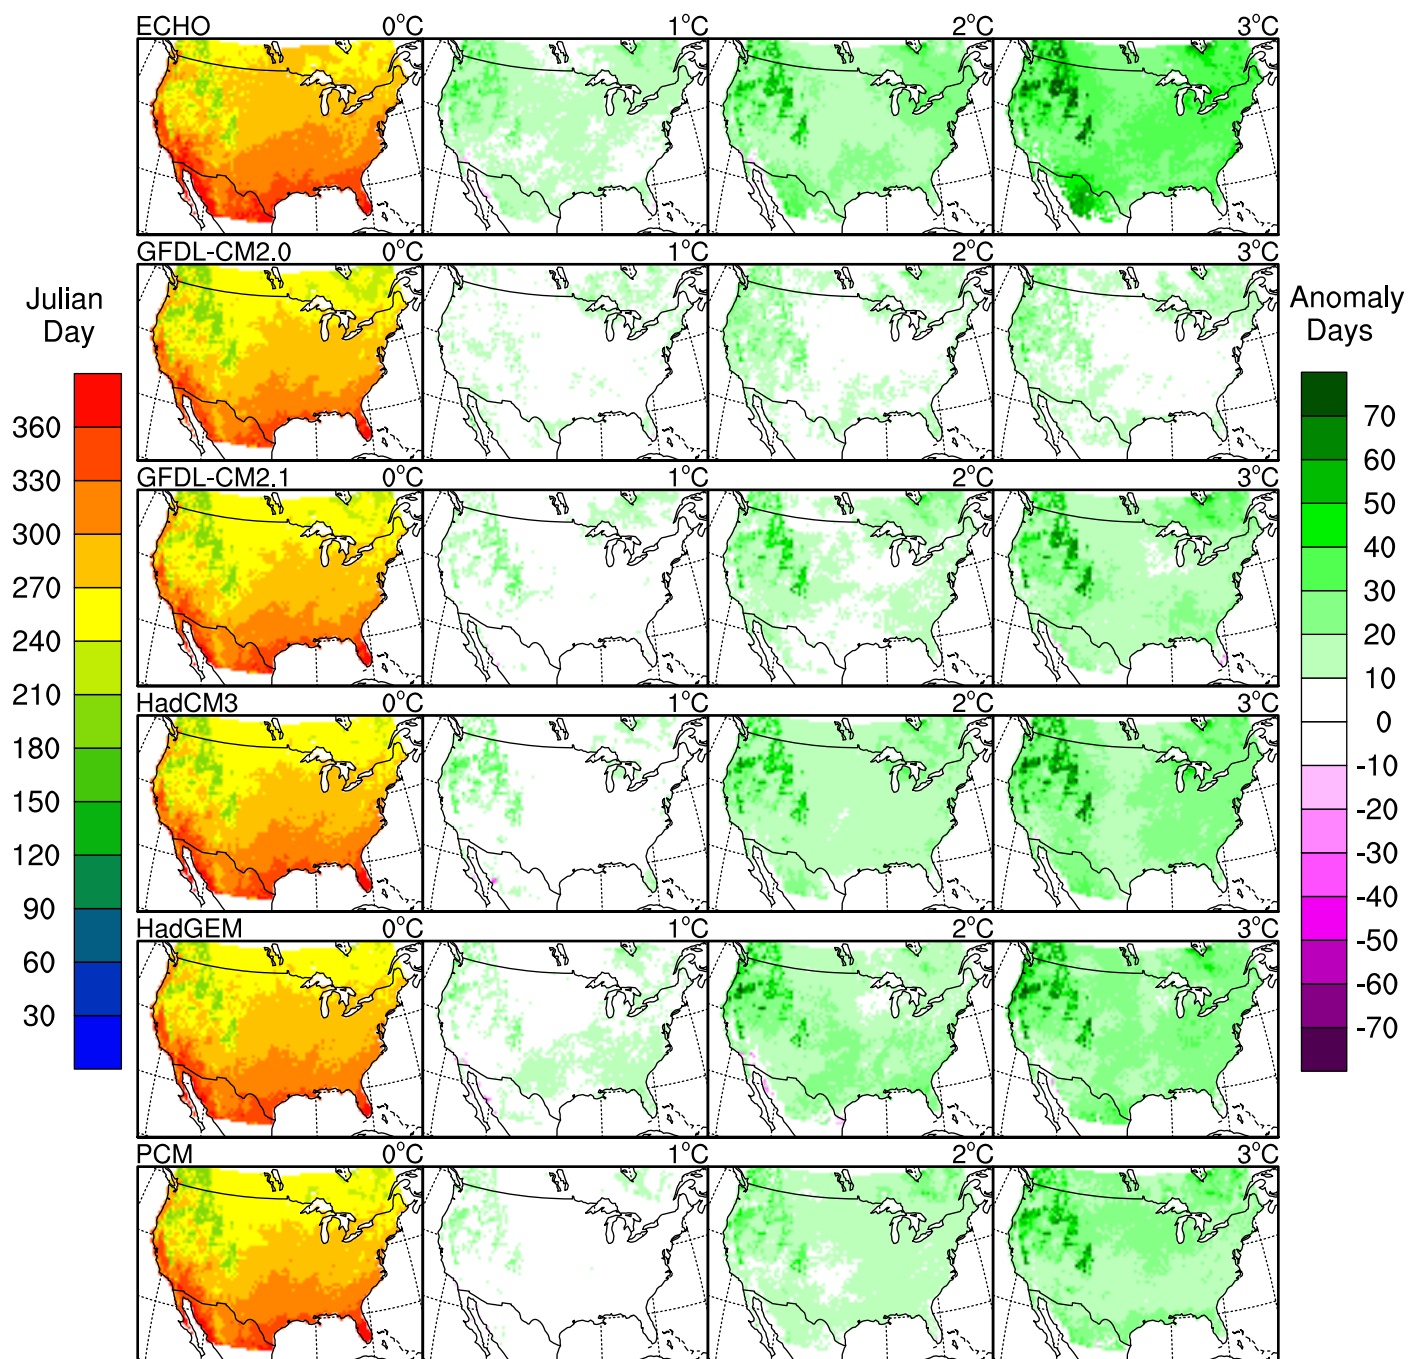
Beginning of Growing Season (Julian Day)



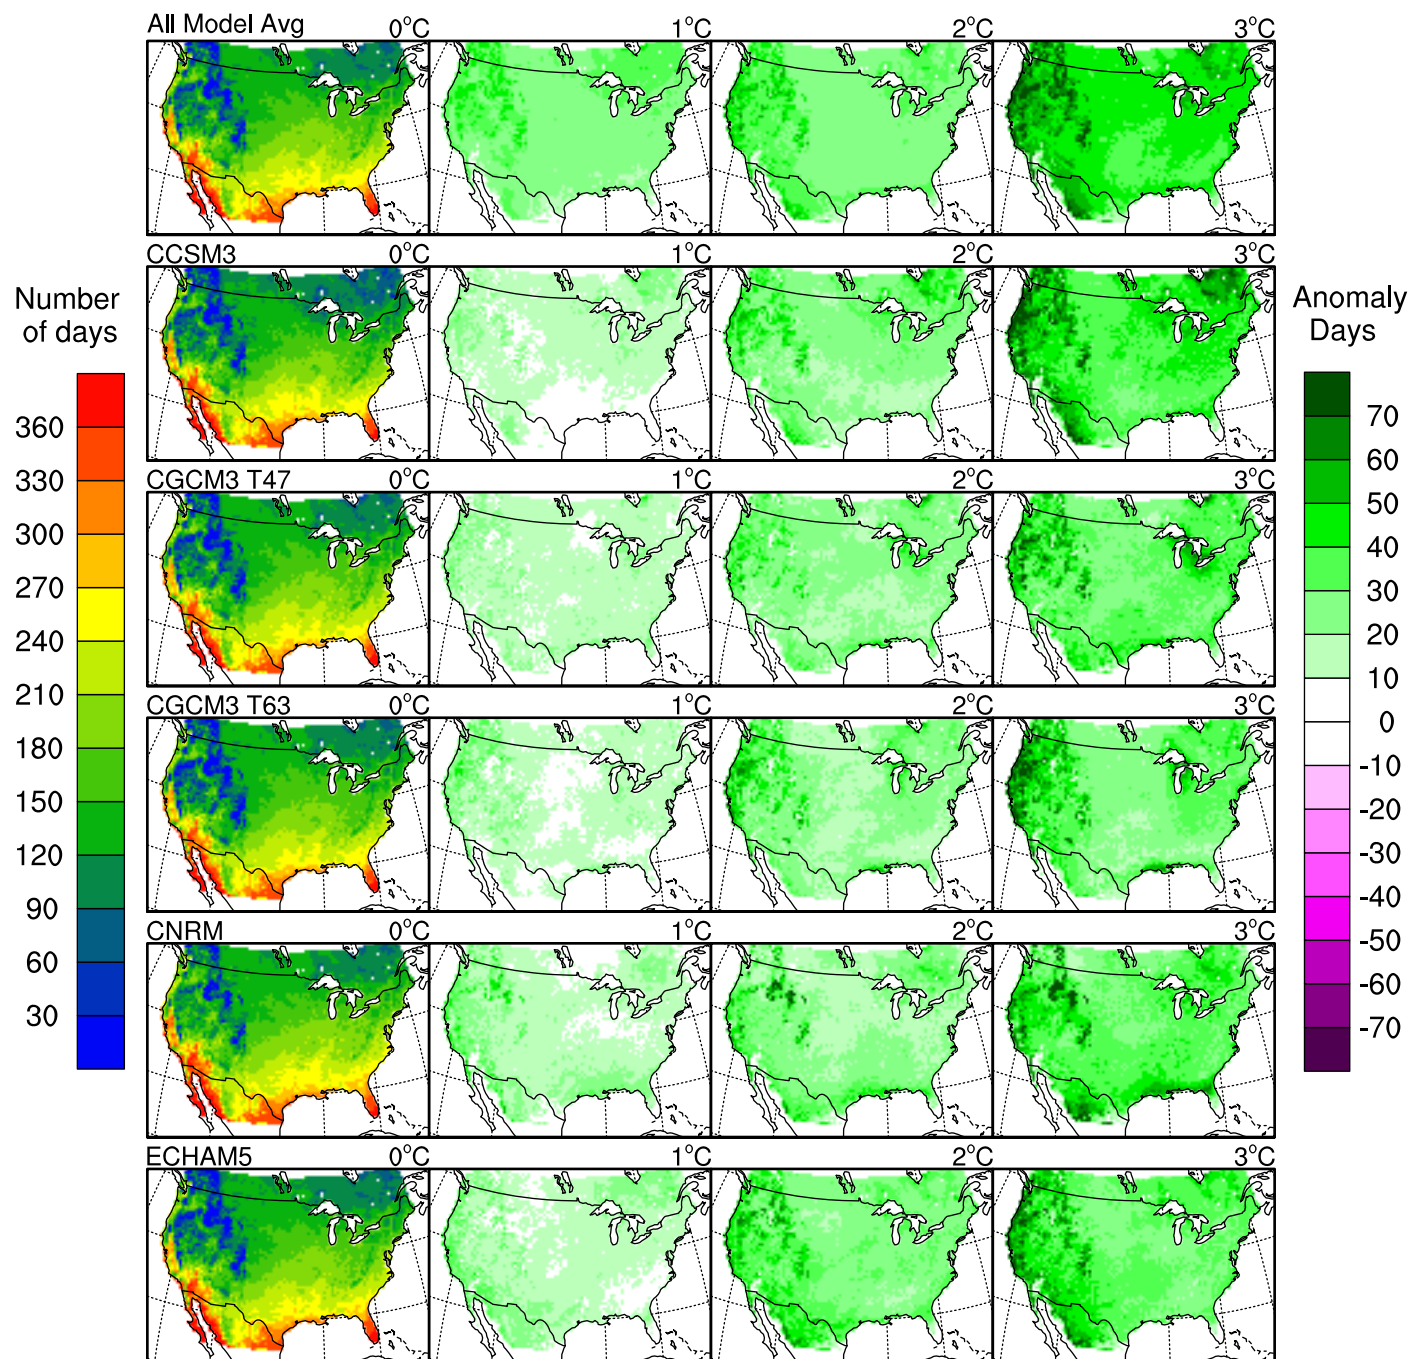
End of Growing Season (Julian Day)



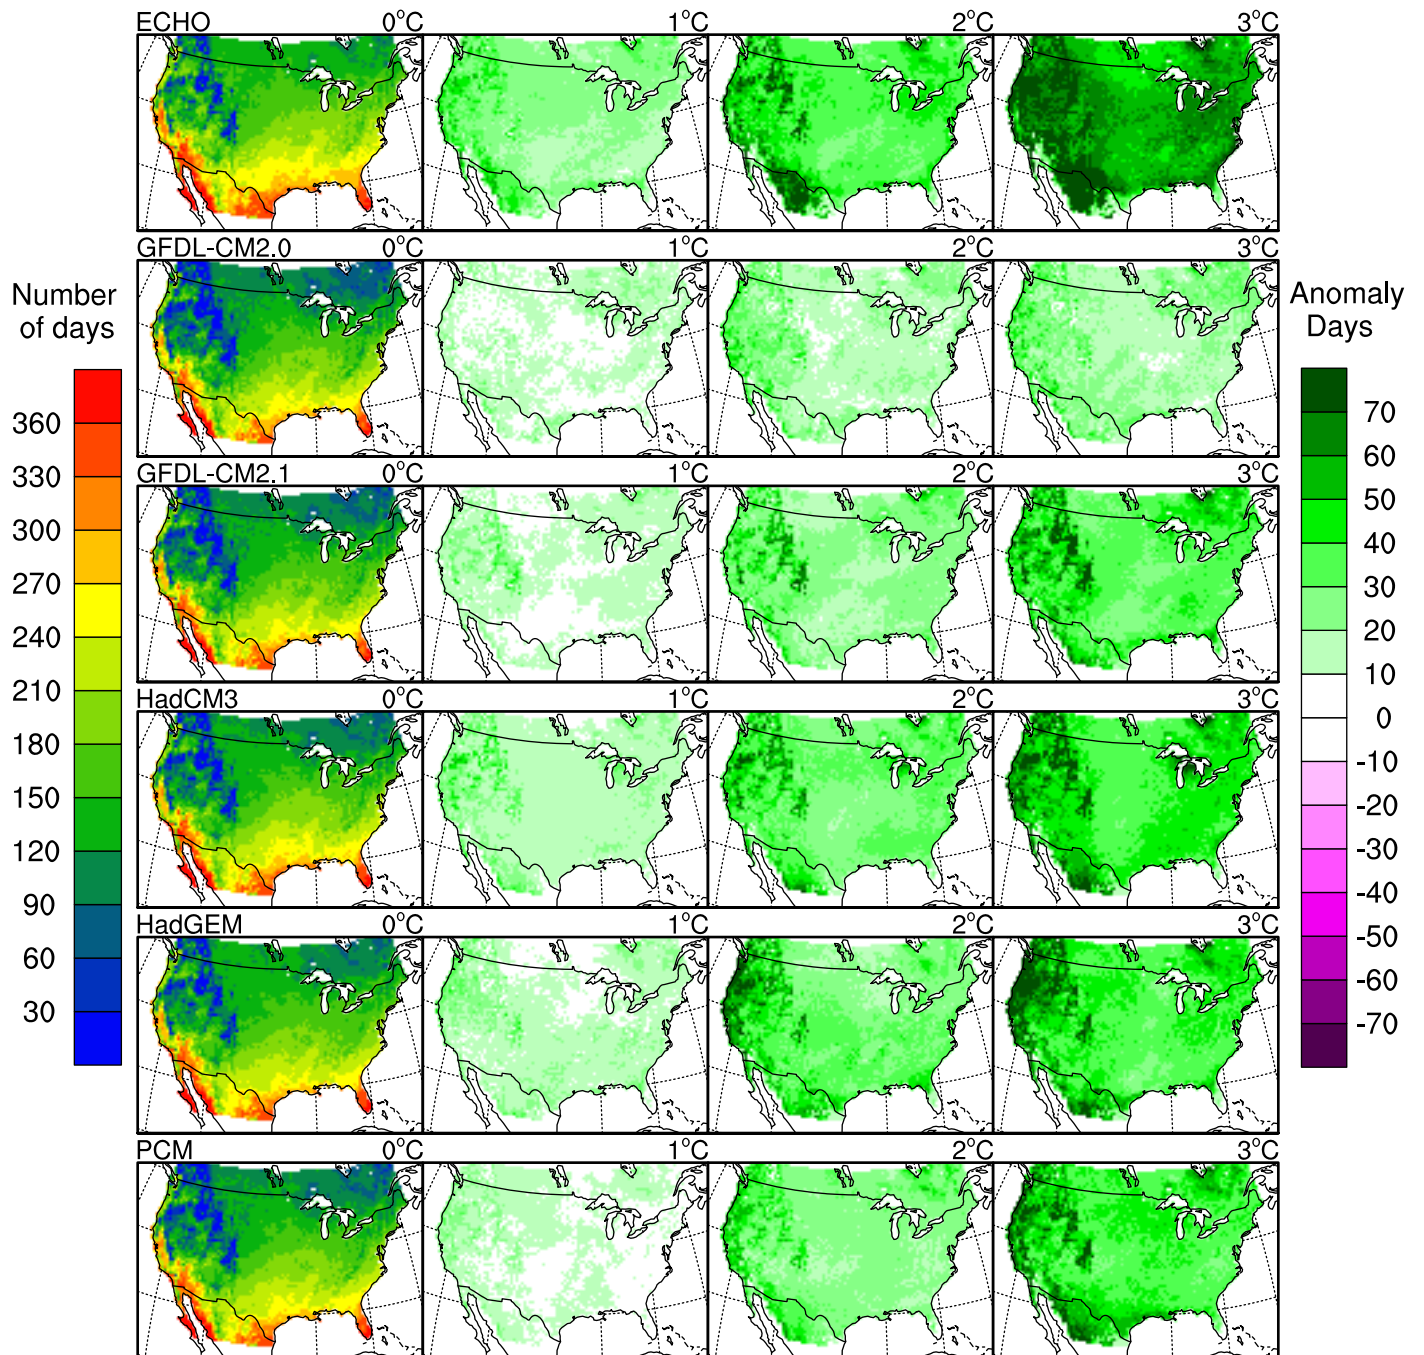
End of Growing Season (Julian Day)



Length of Growing Season (Julian Day)



Length of Growing Season (Julian Day)



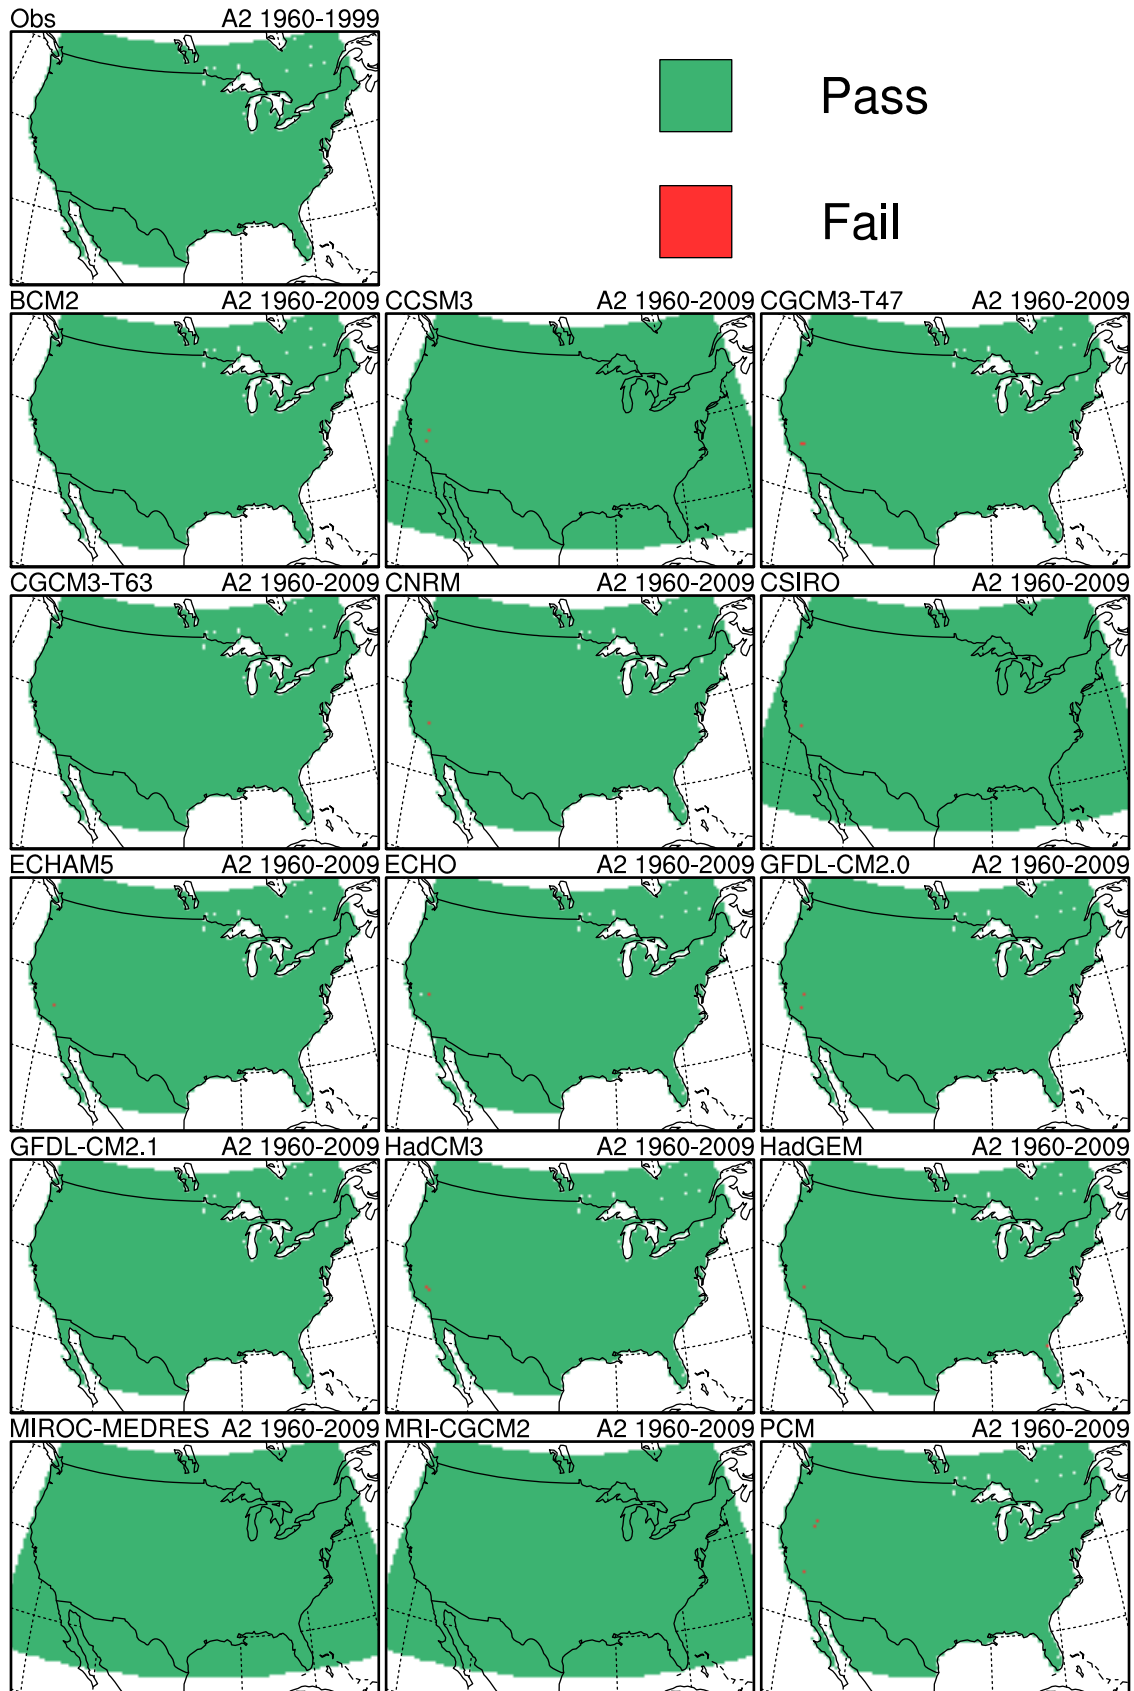
APPENDIX D. Error Plots

This Appendix provides maps of temperature and precipitation error tests applied to the gridded observations and gridded downscaling. The tests are described in Table 5. Some error tests are pass/fail; these are indicated by green (pass) and red (fail) scales. Other tests allow for multiple results; these use a color scale.

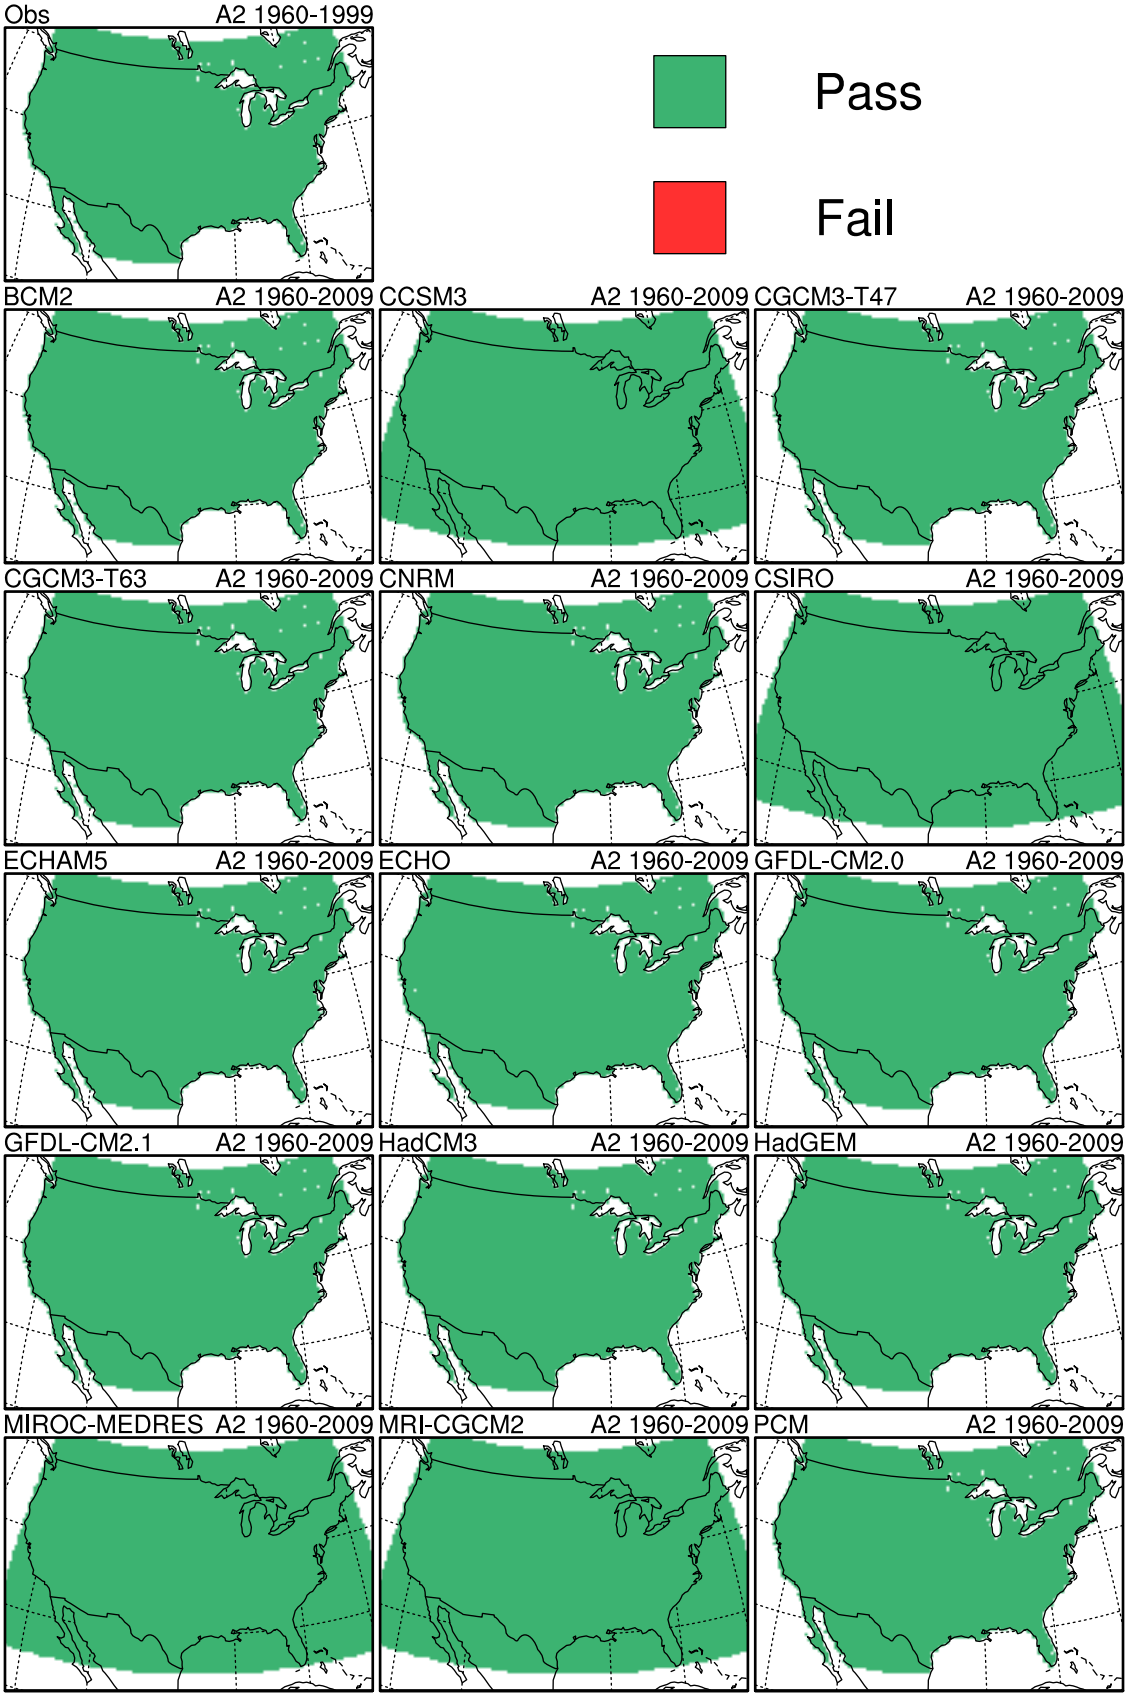
Thanks to Adam Terando (NCSU/USGS) for suggestions on additional error tests to include.

Original postscript files are also provided in electronic format.

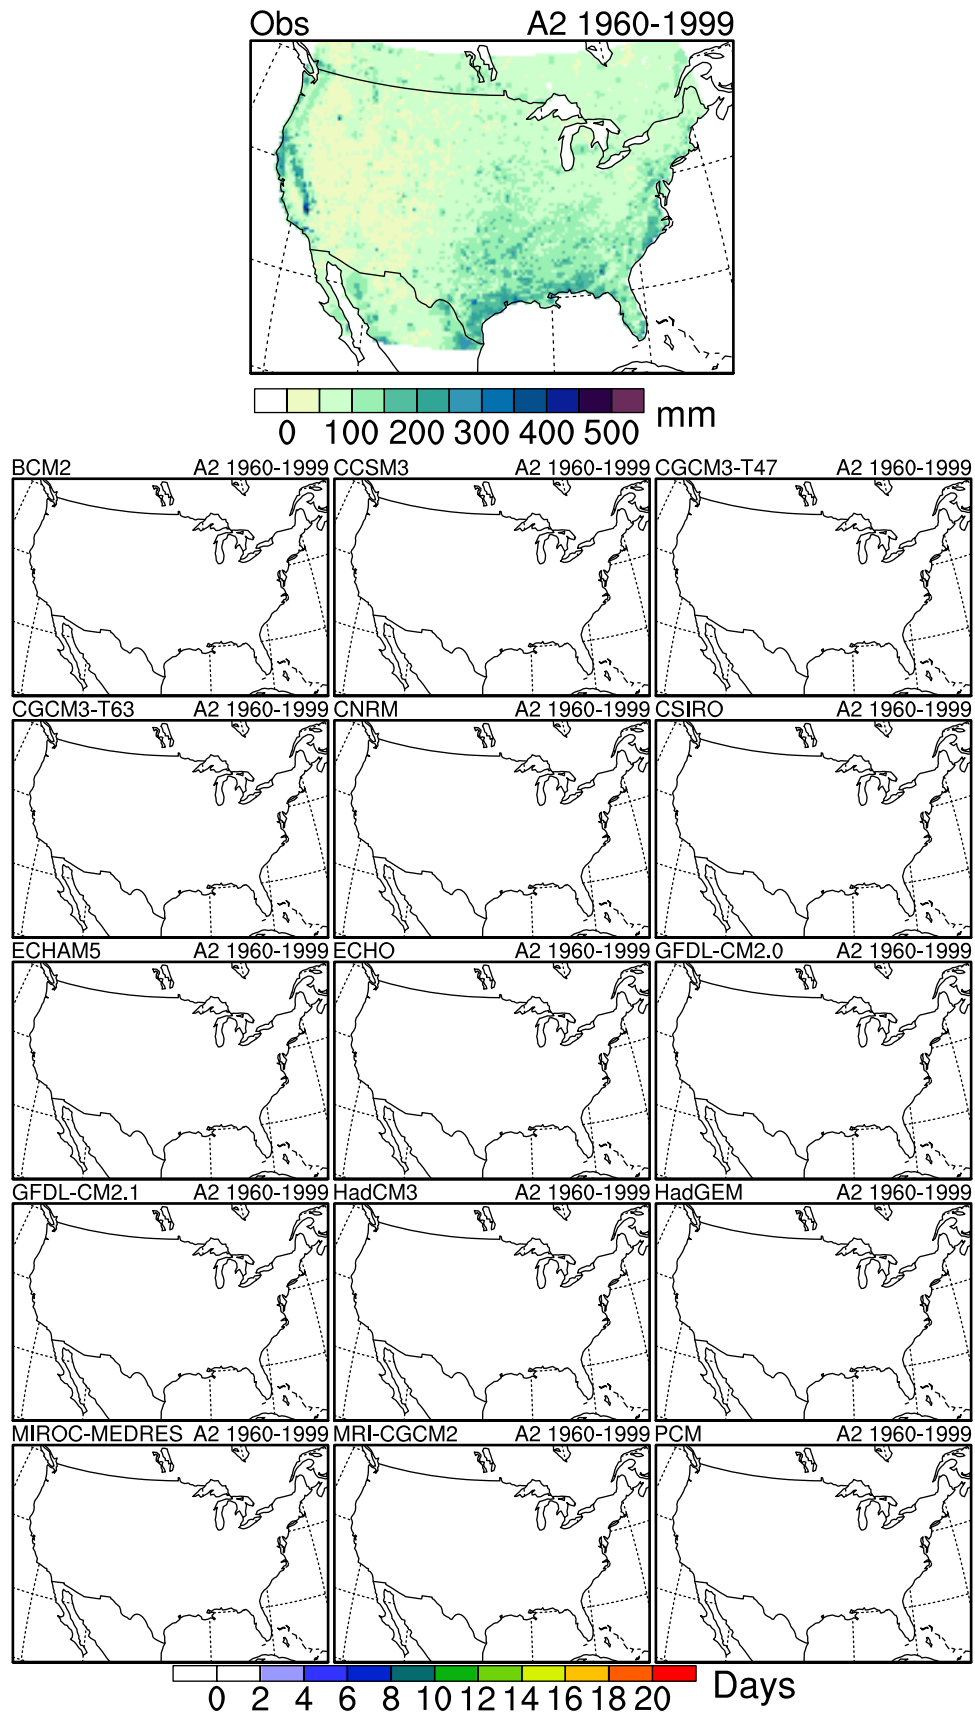
Occurrences of $\max(\text{pr}) > \max$ observed CONUS value



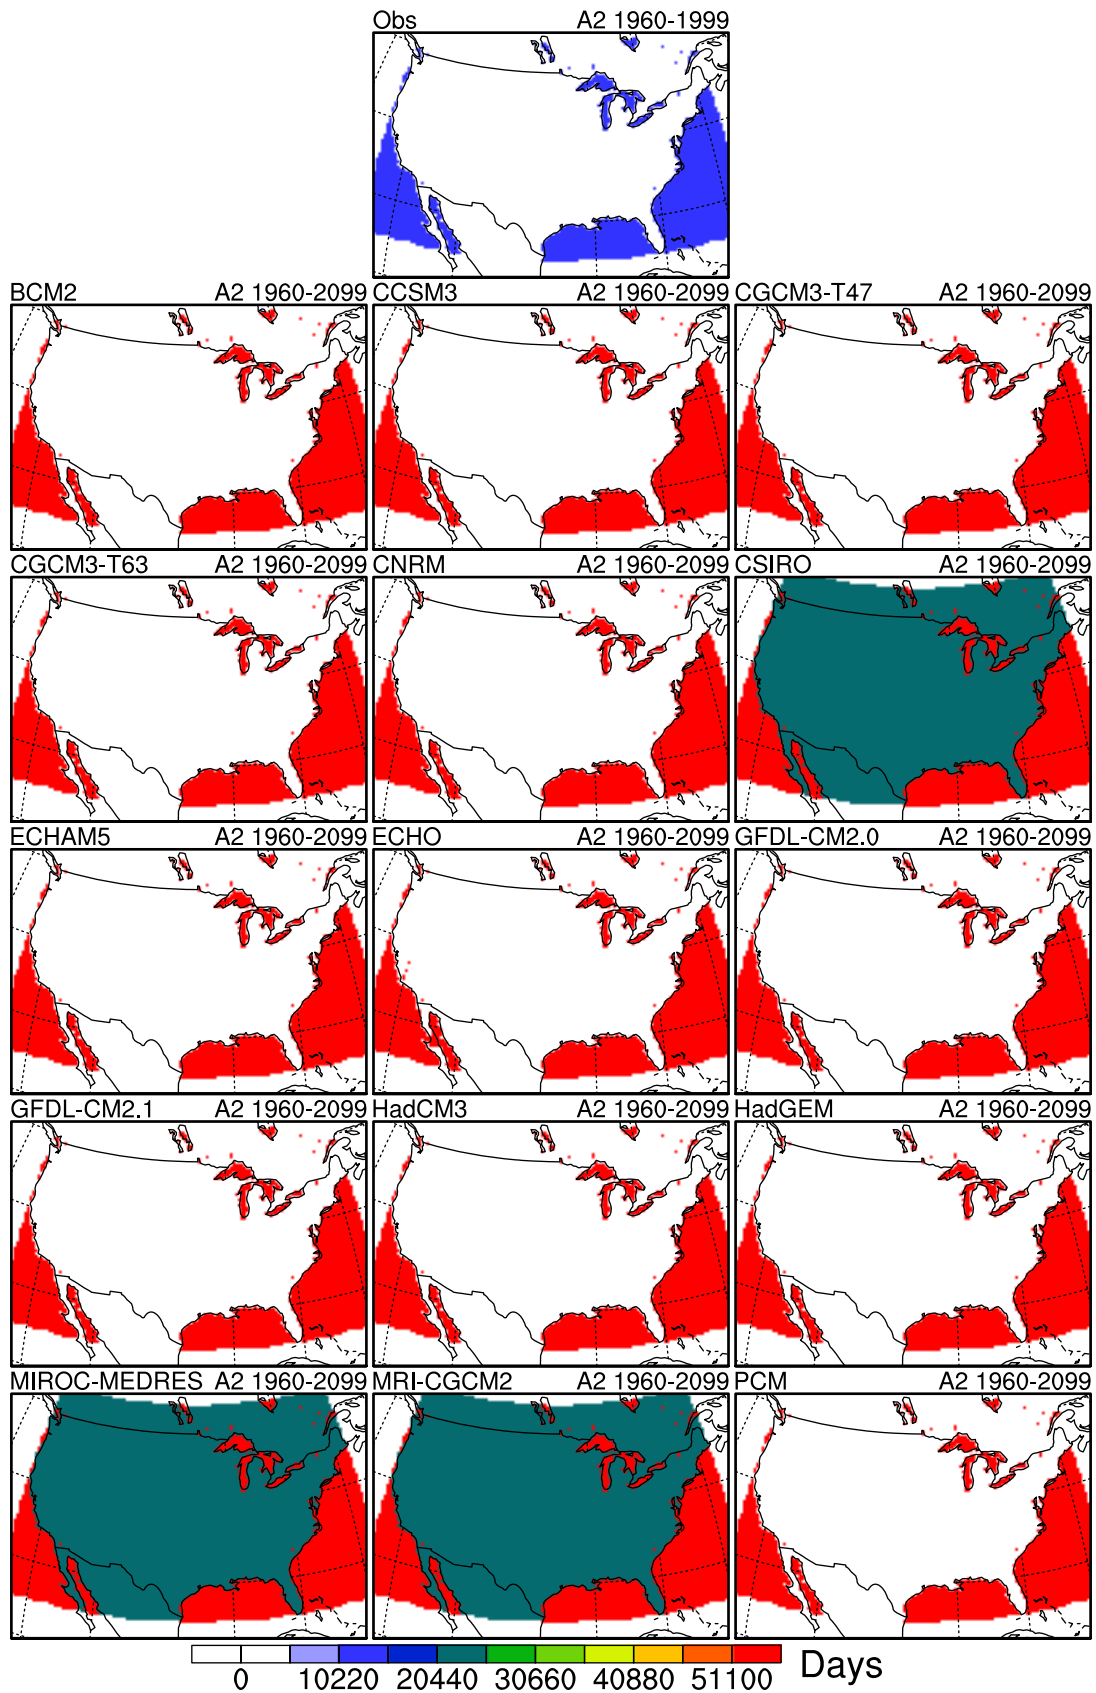
Occurrences of $pr < 0$



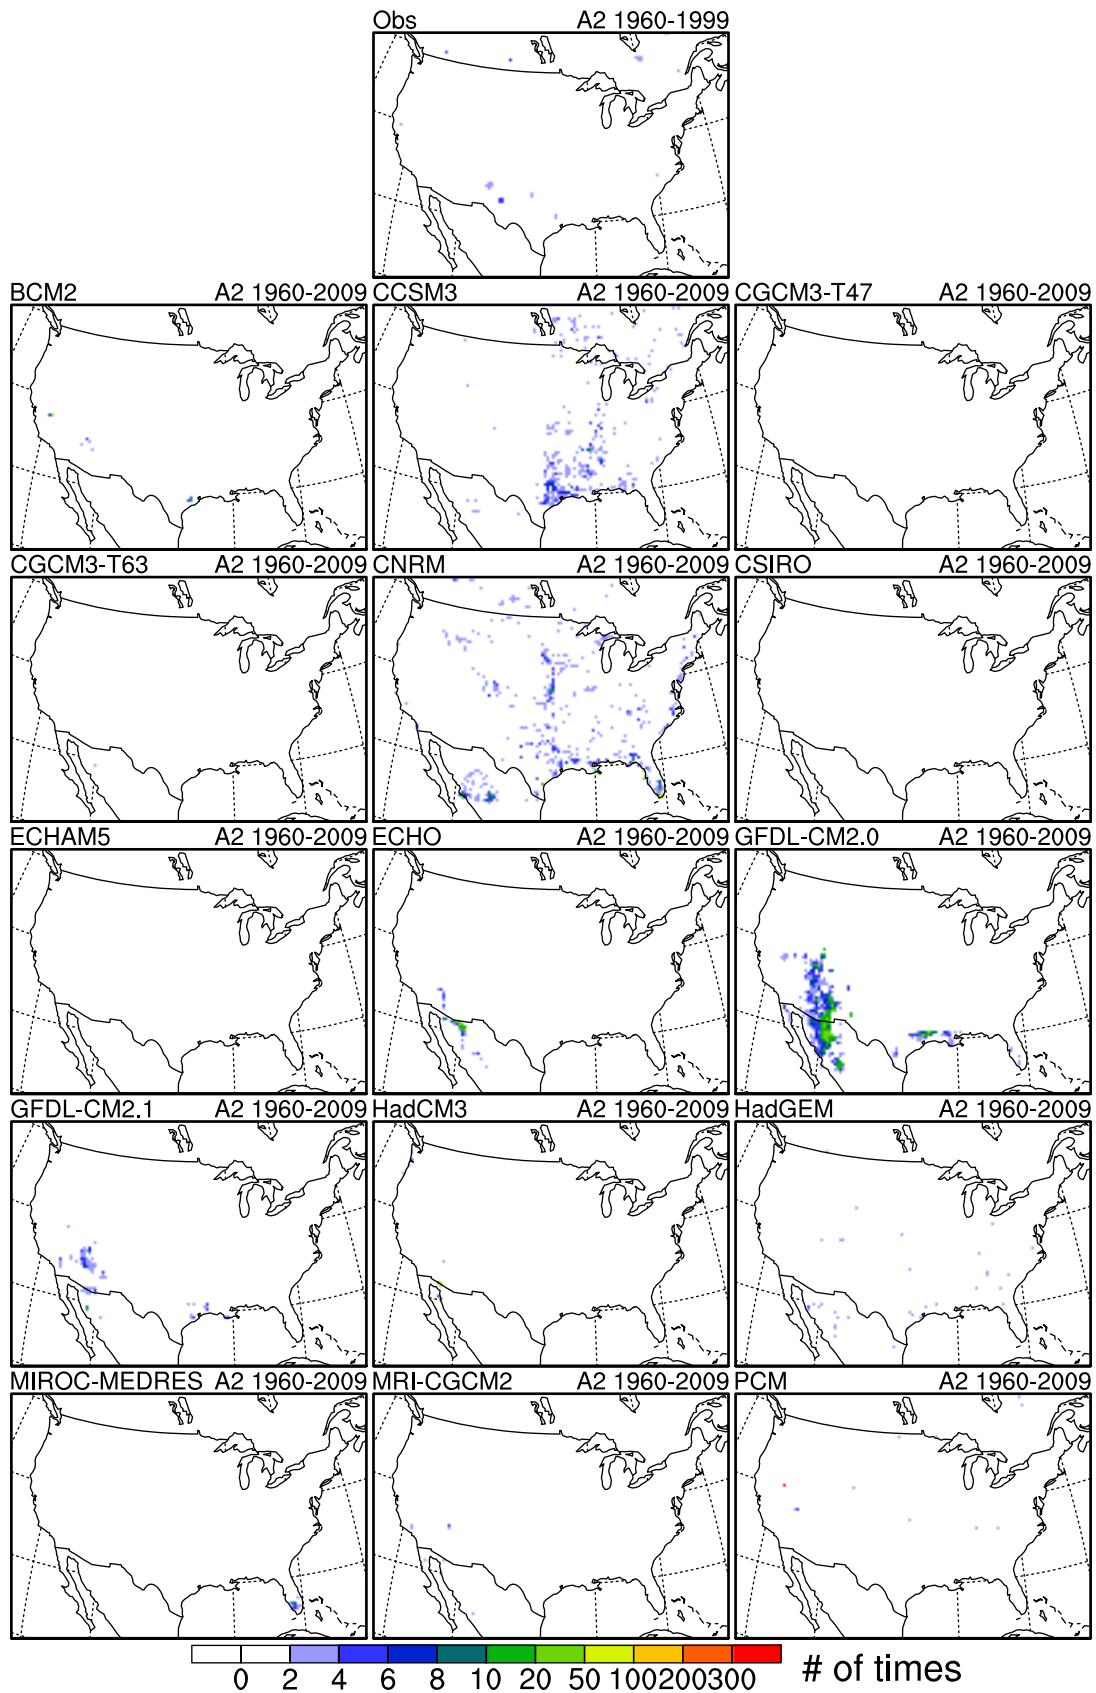
Number of historical days with $pr > \max$ observed local pr value



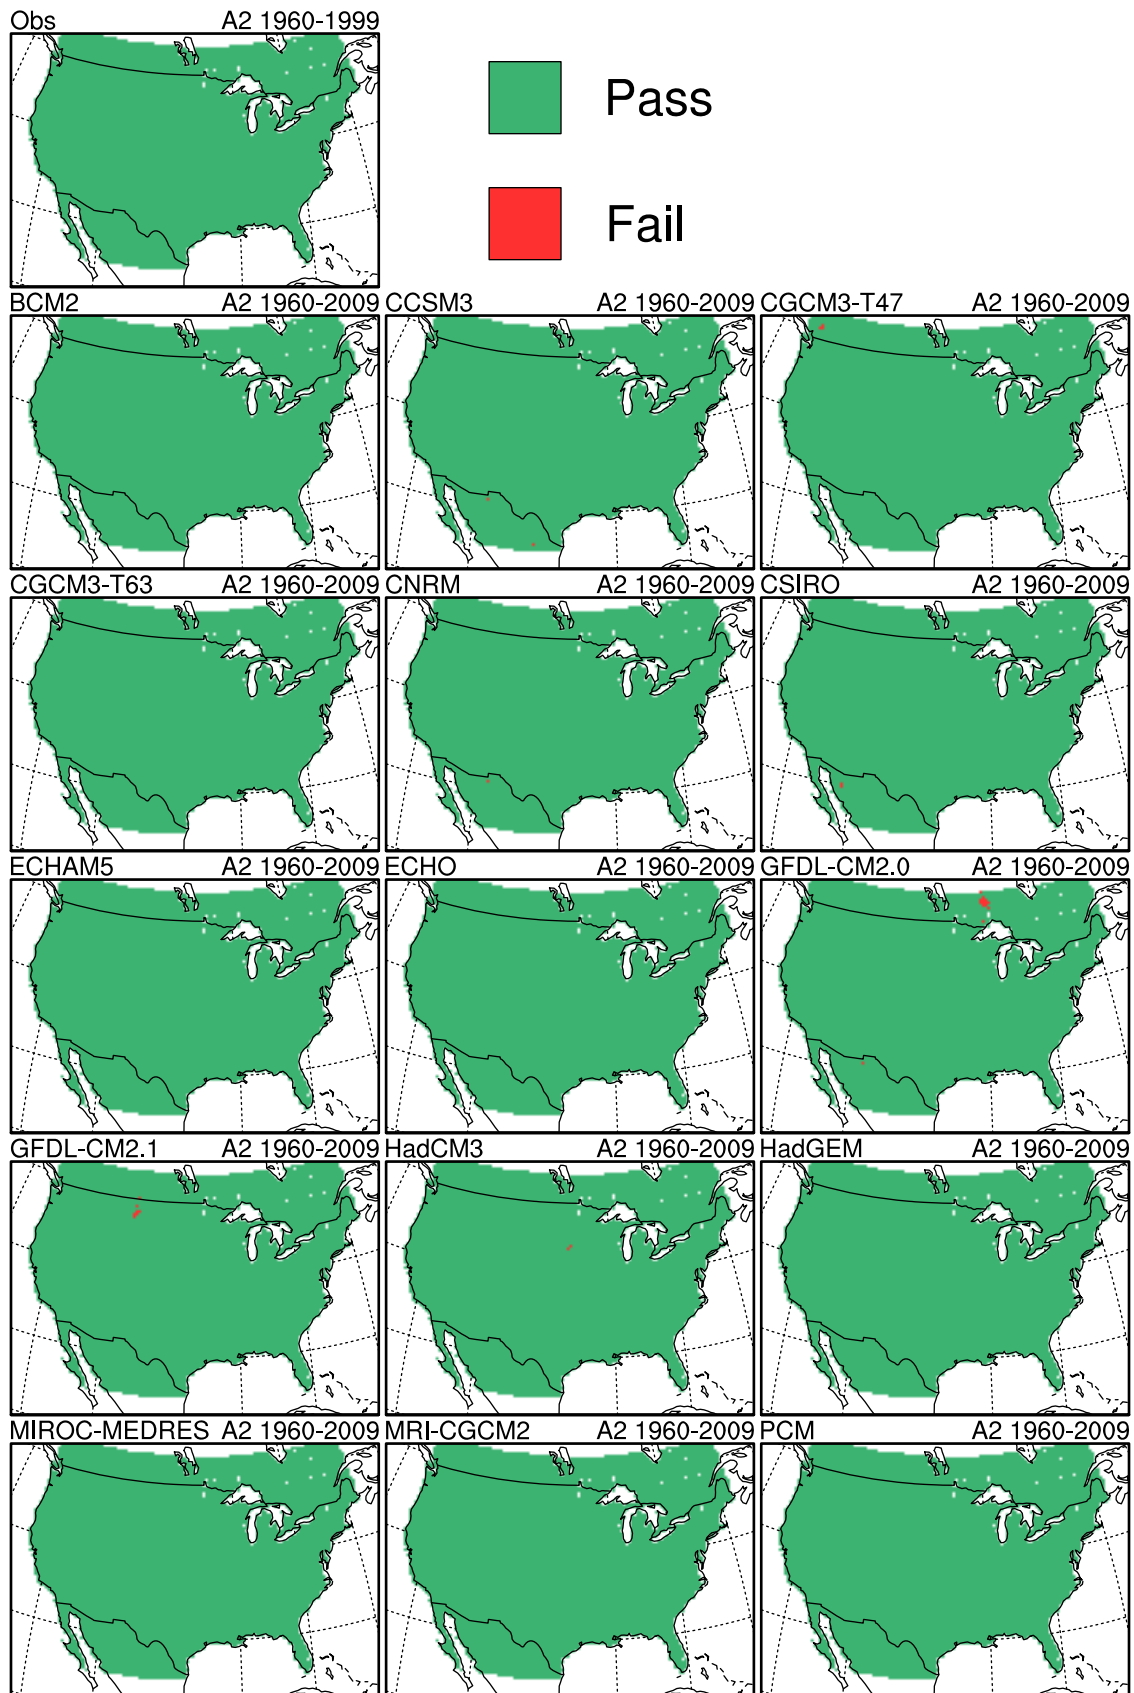
Number of missing pr values



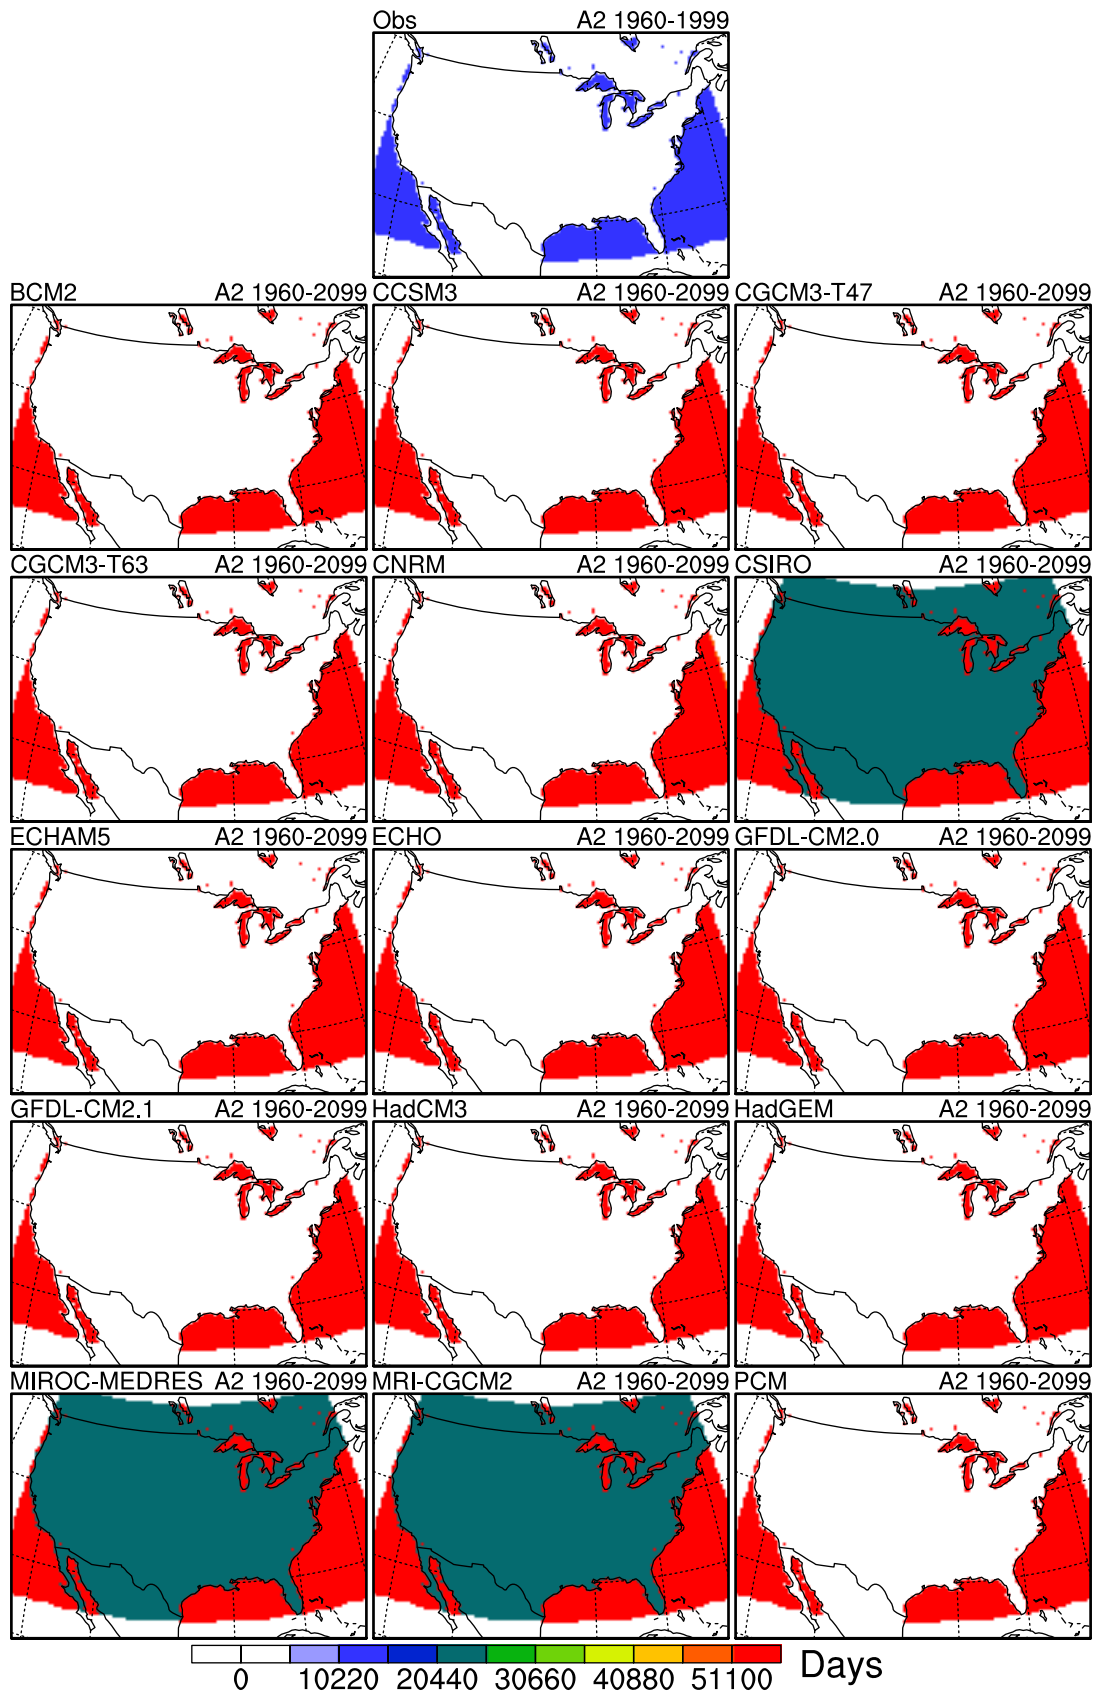
Number of times a pr value repeats > 3x



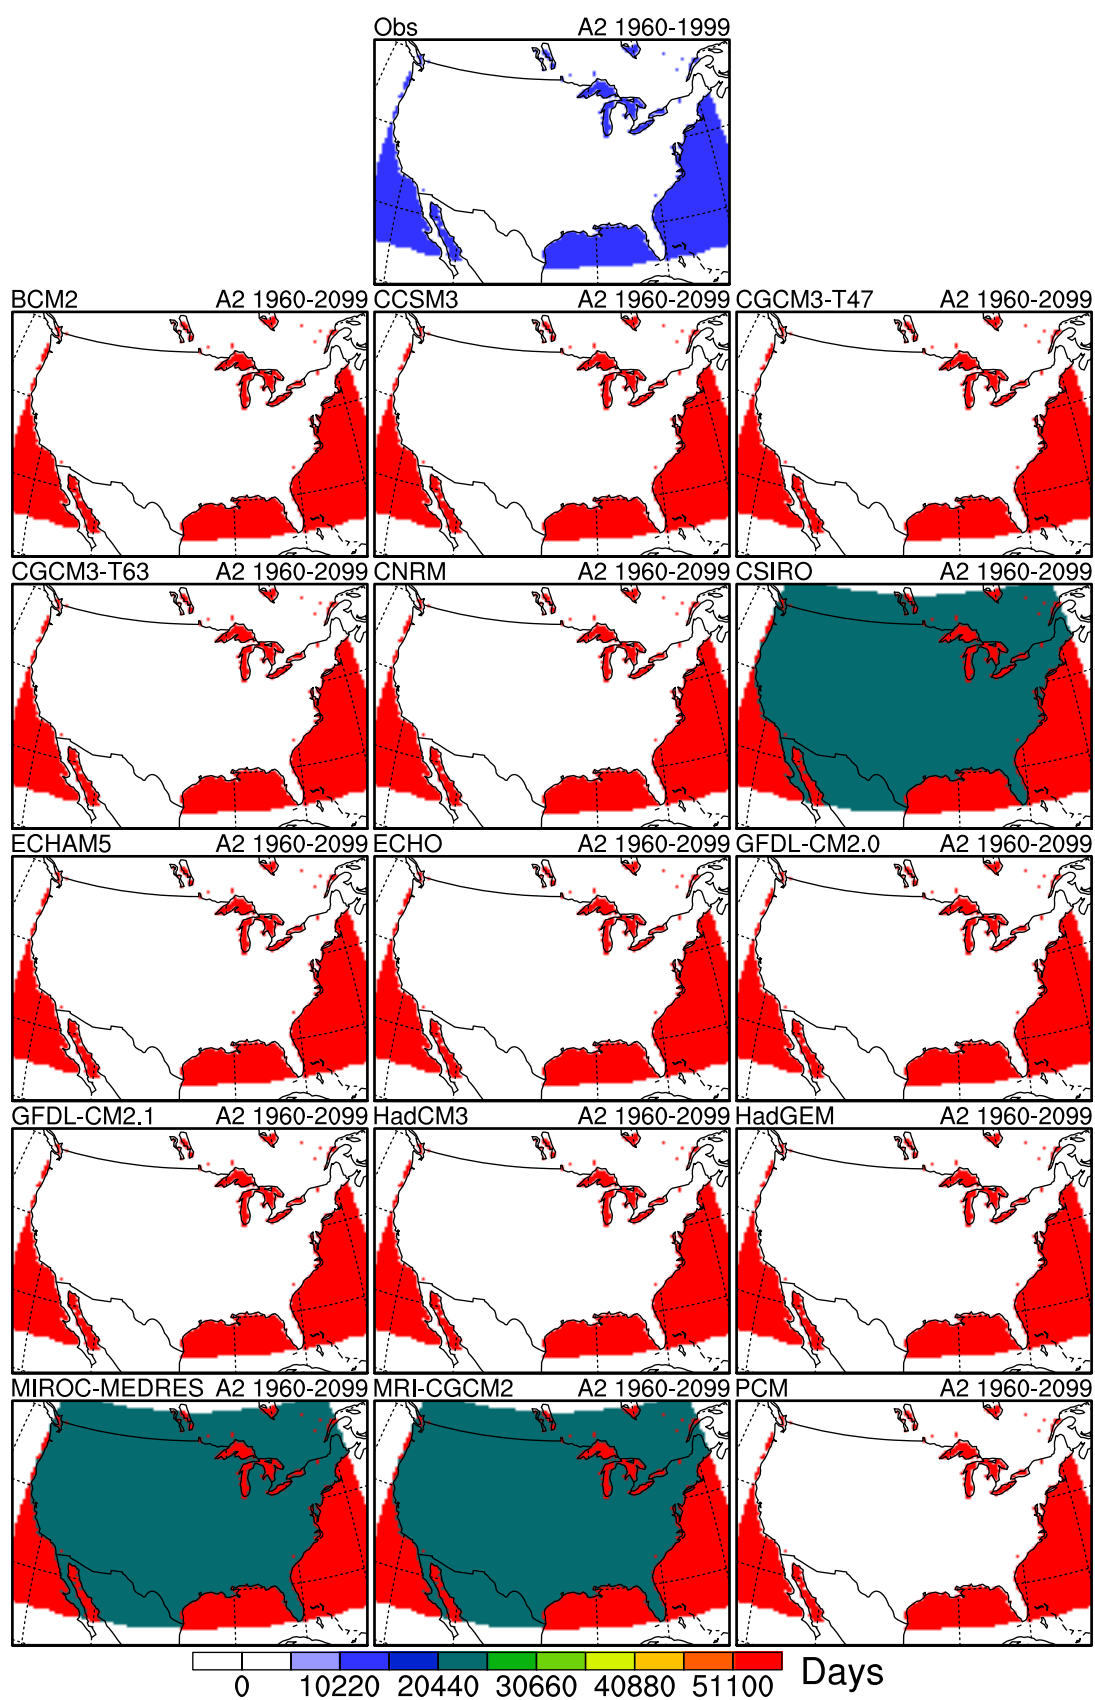
Occurrences of $\max(t_{\max}) > \max$ observed CONUS value



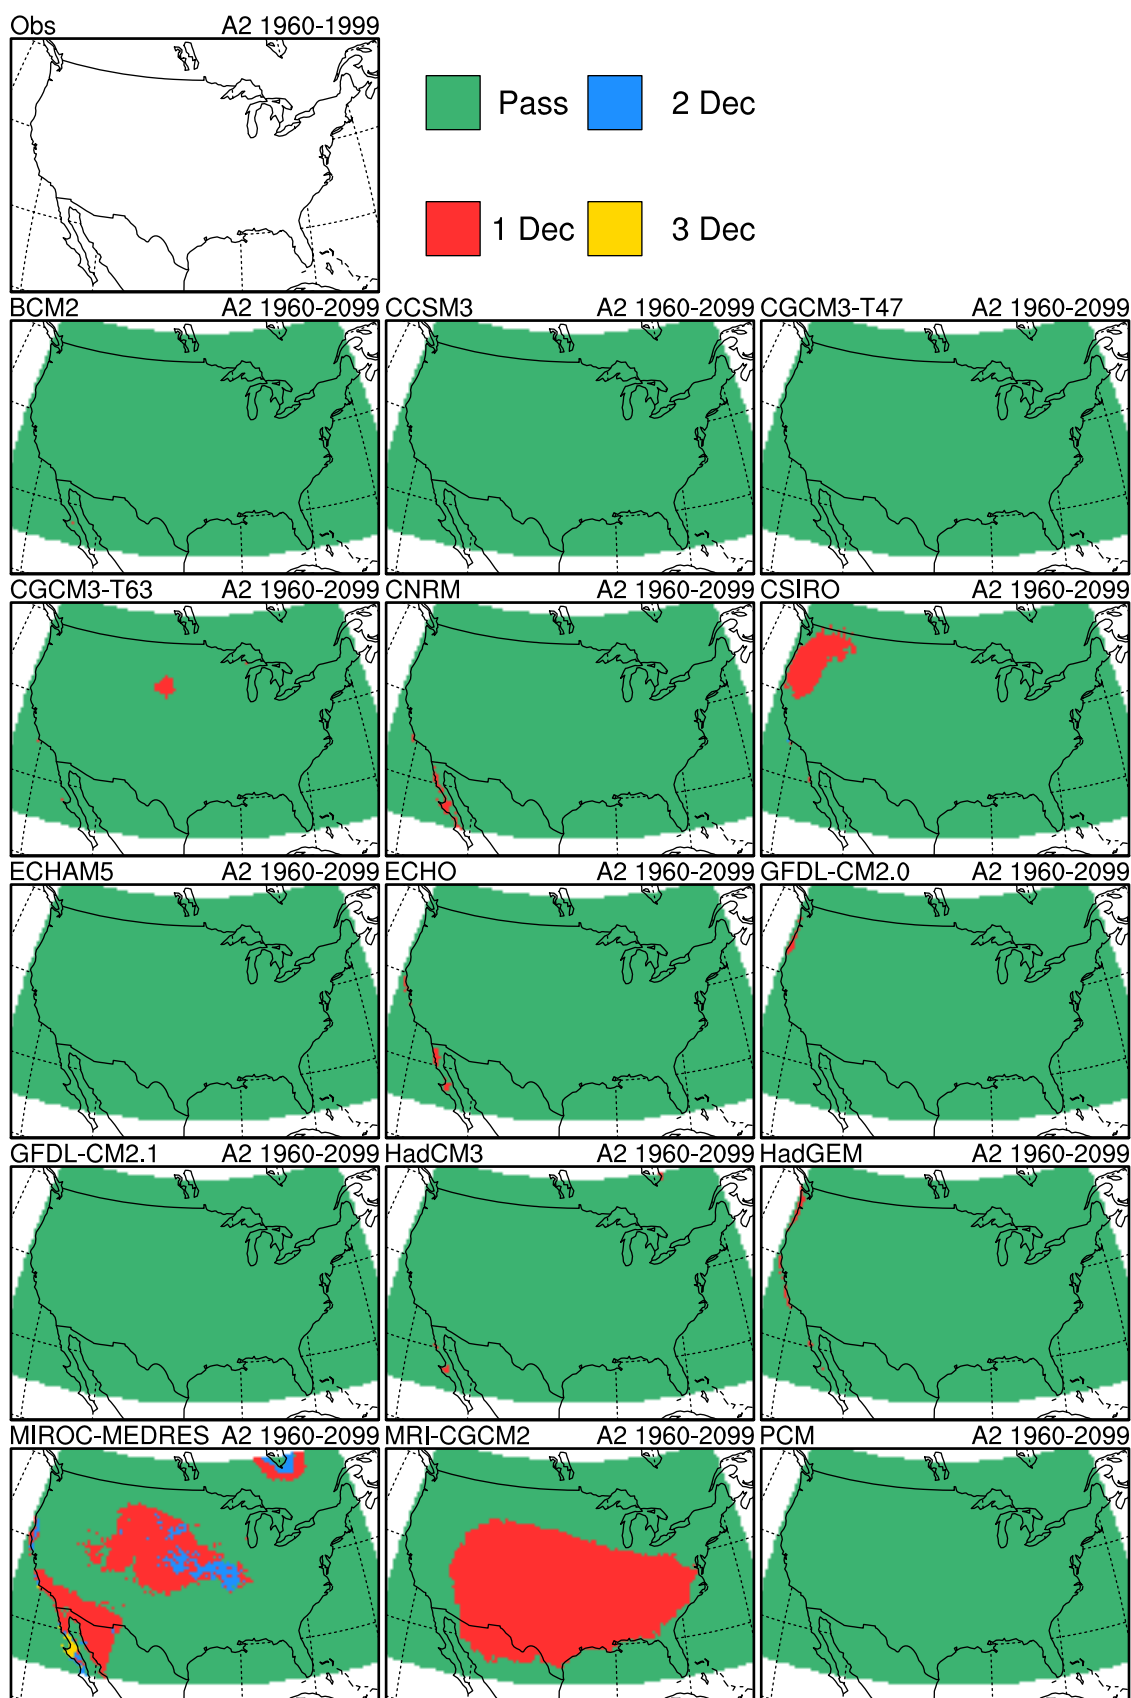
Number of missing tmax values



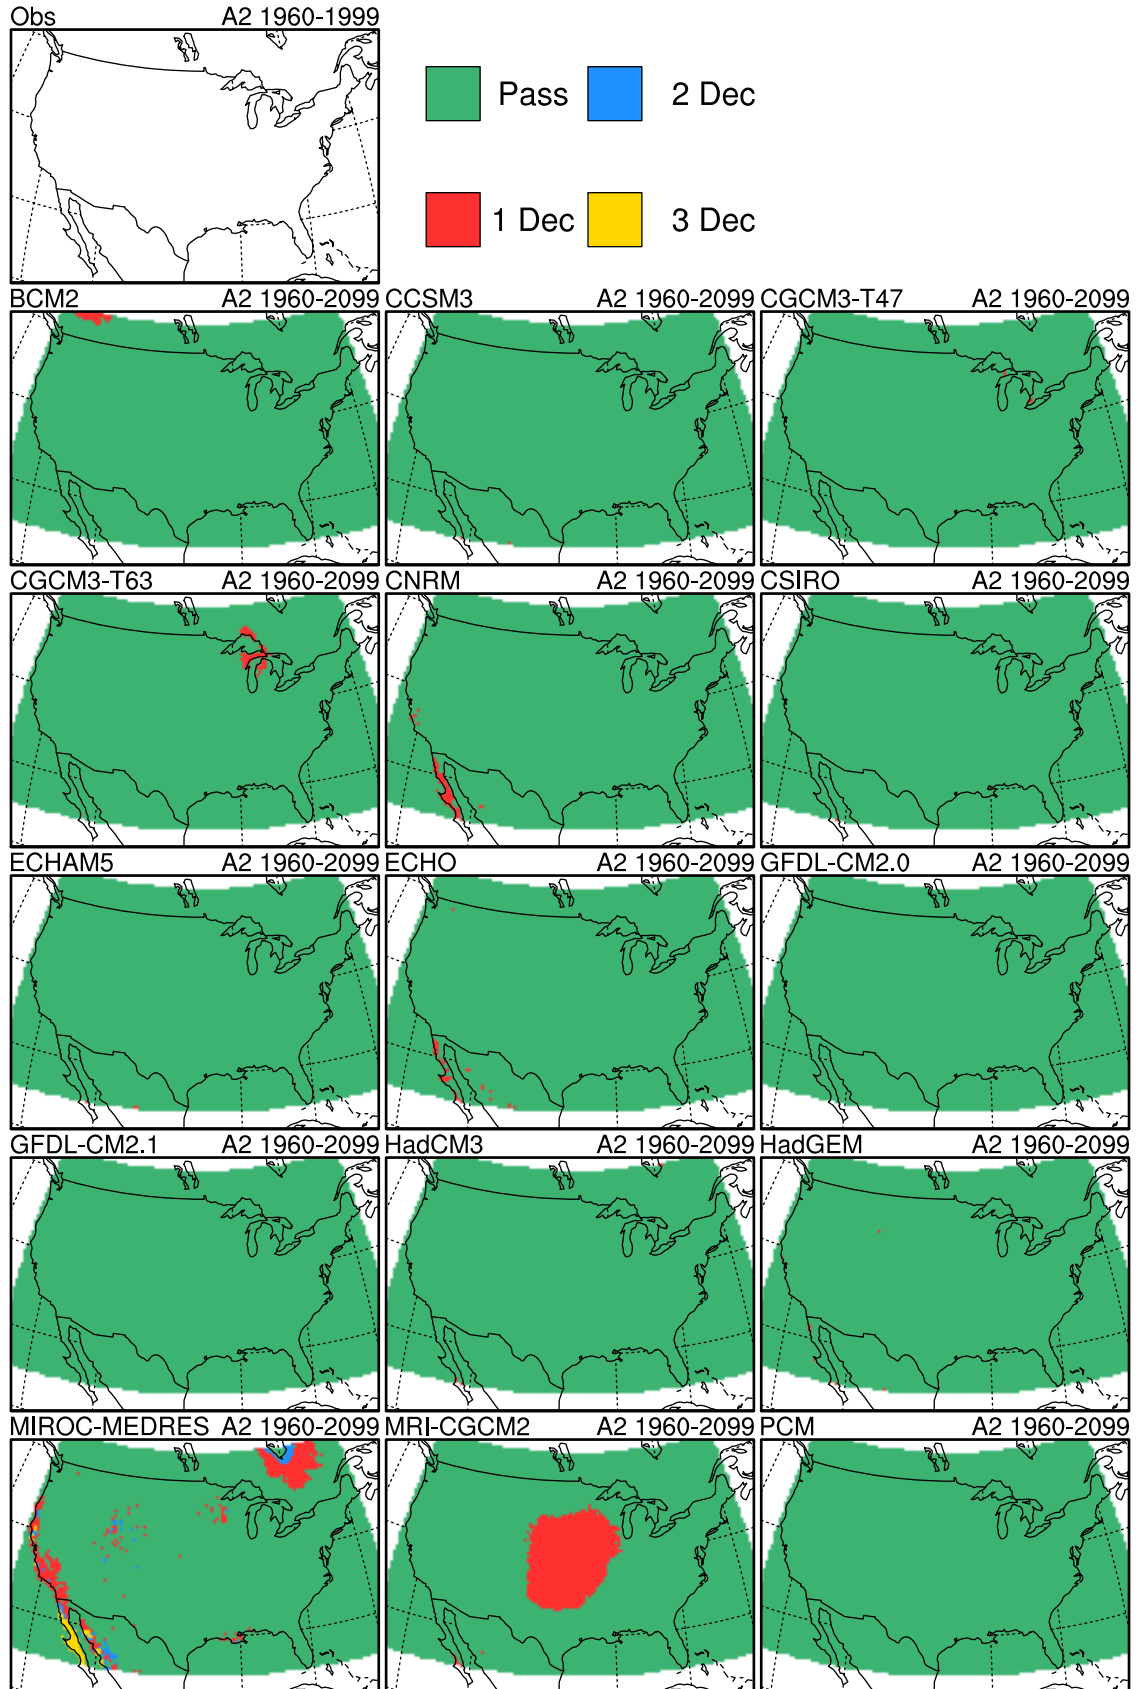
Number of missing tmin values



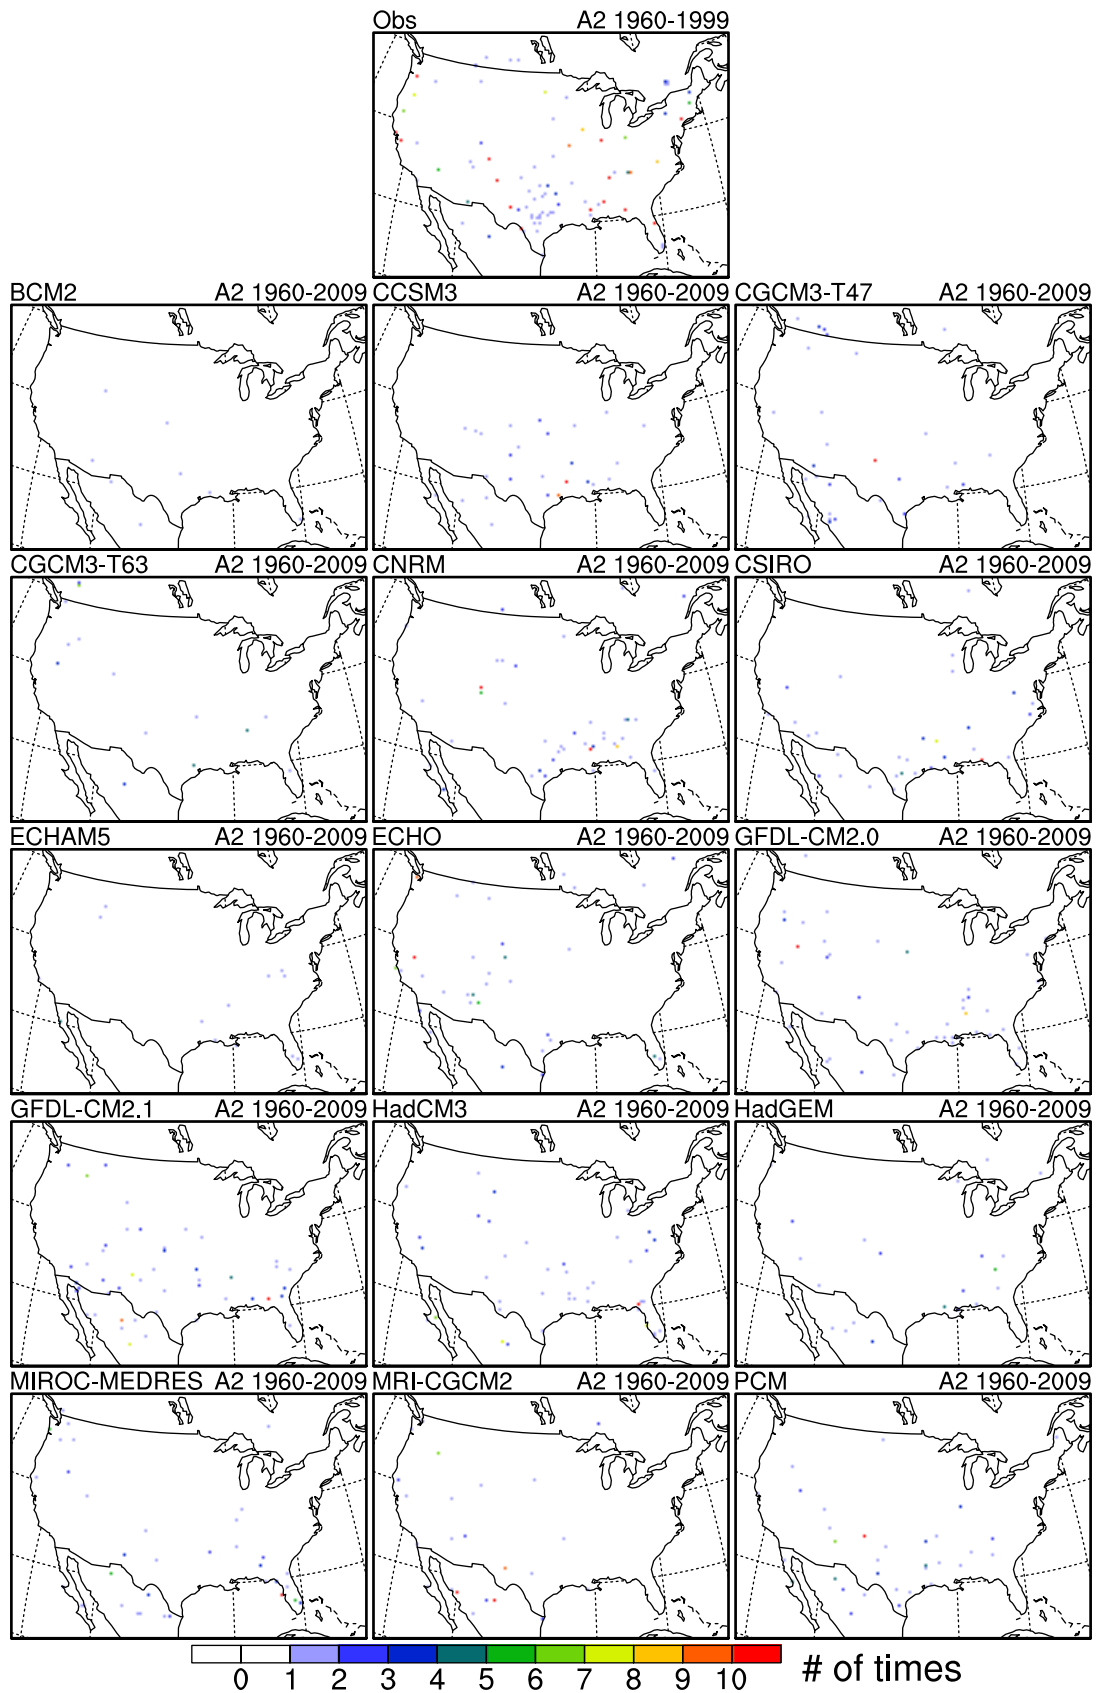
Tmax increases over decadal time scales



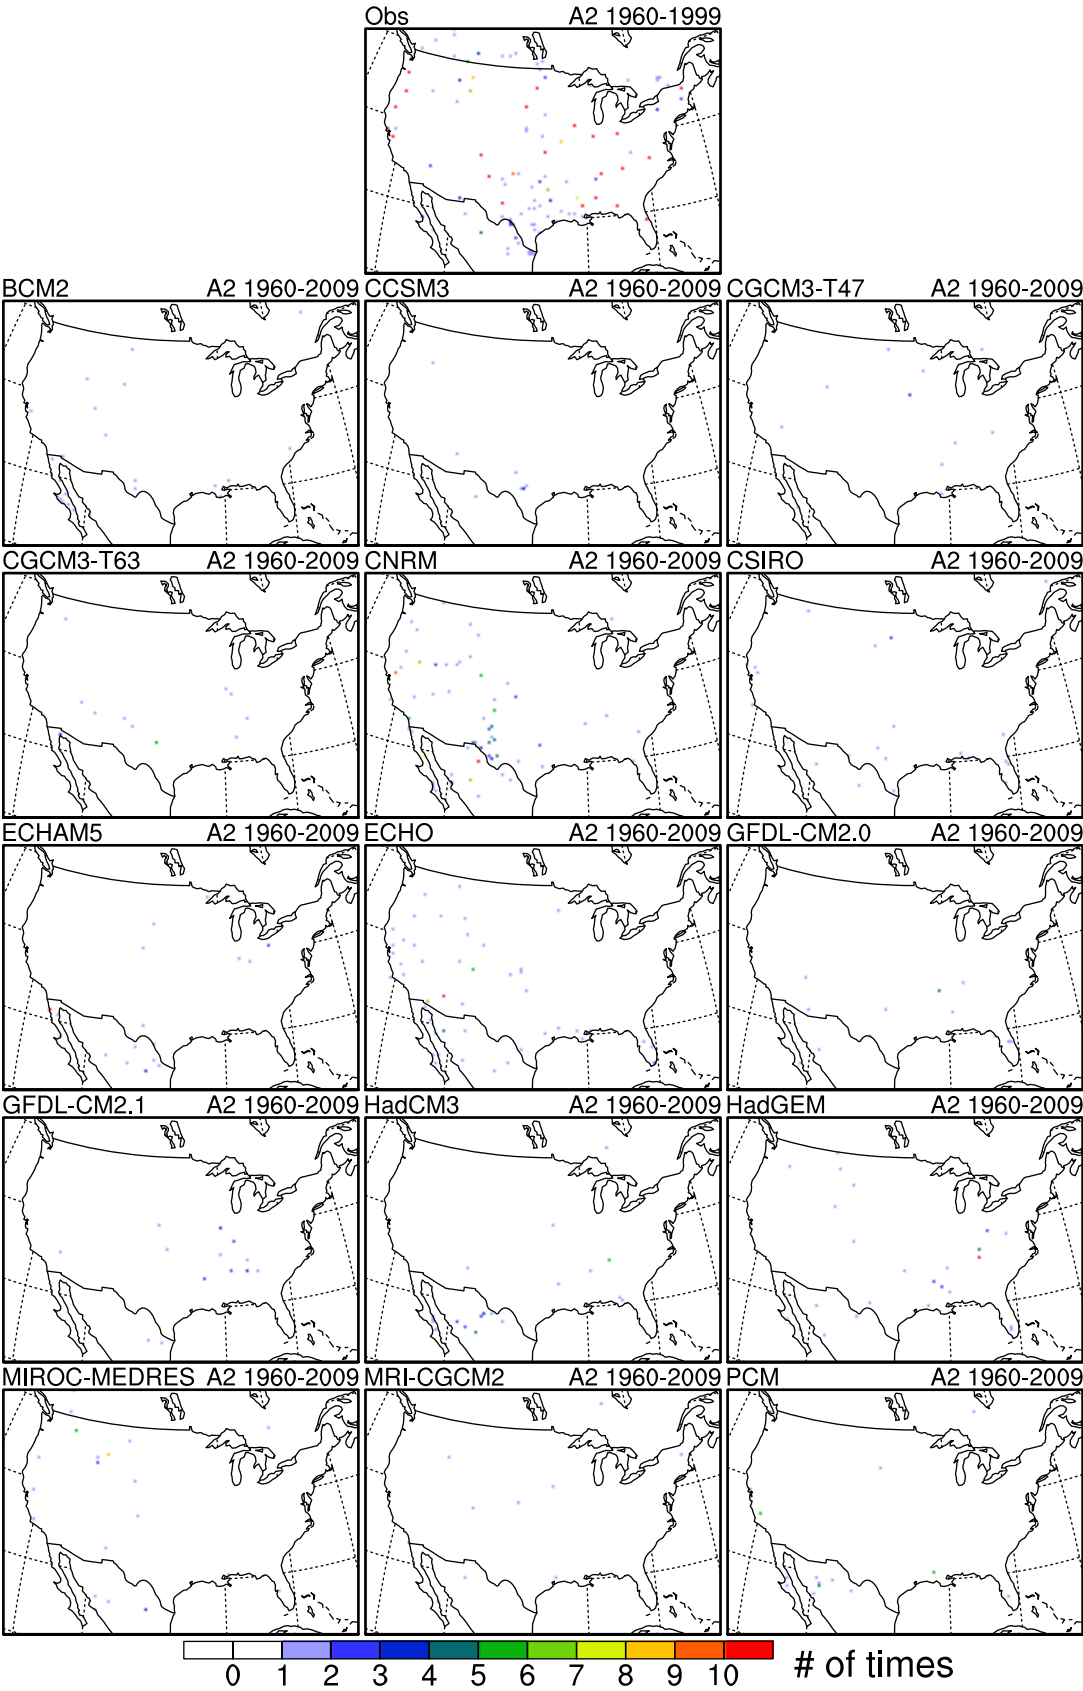
Tmin increases over decadal time scales



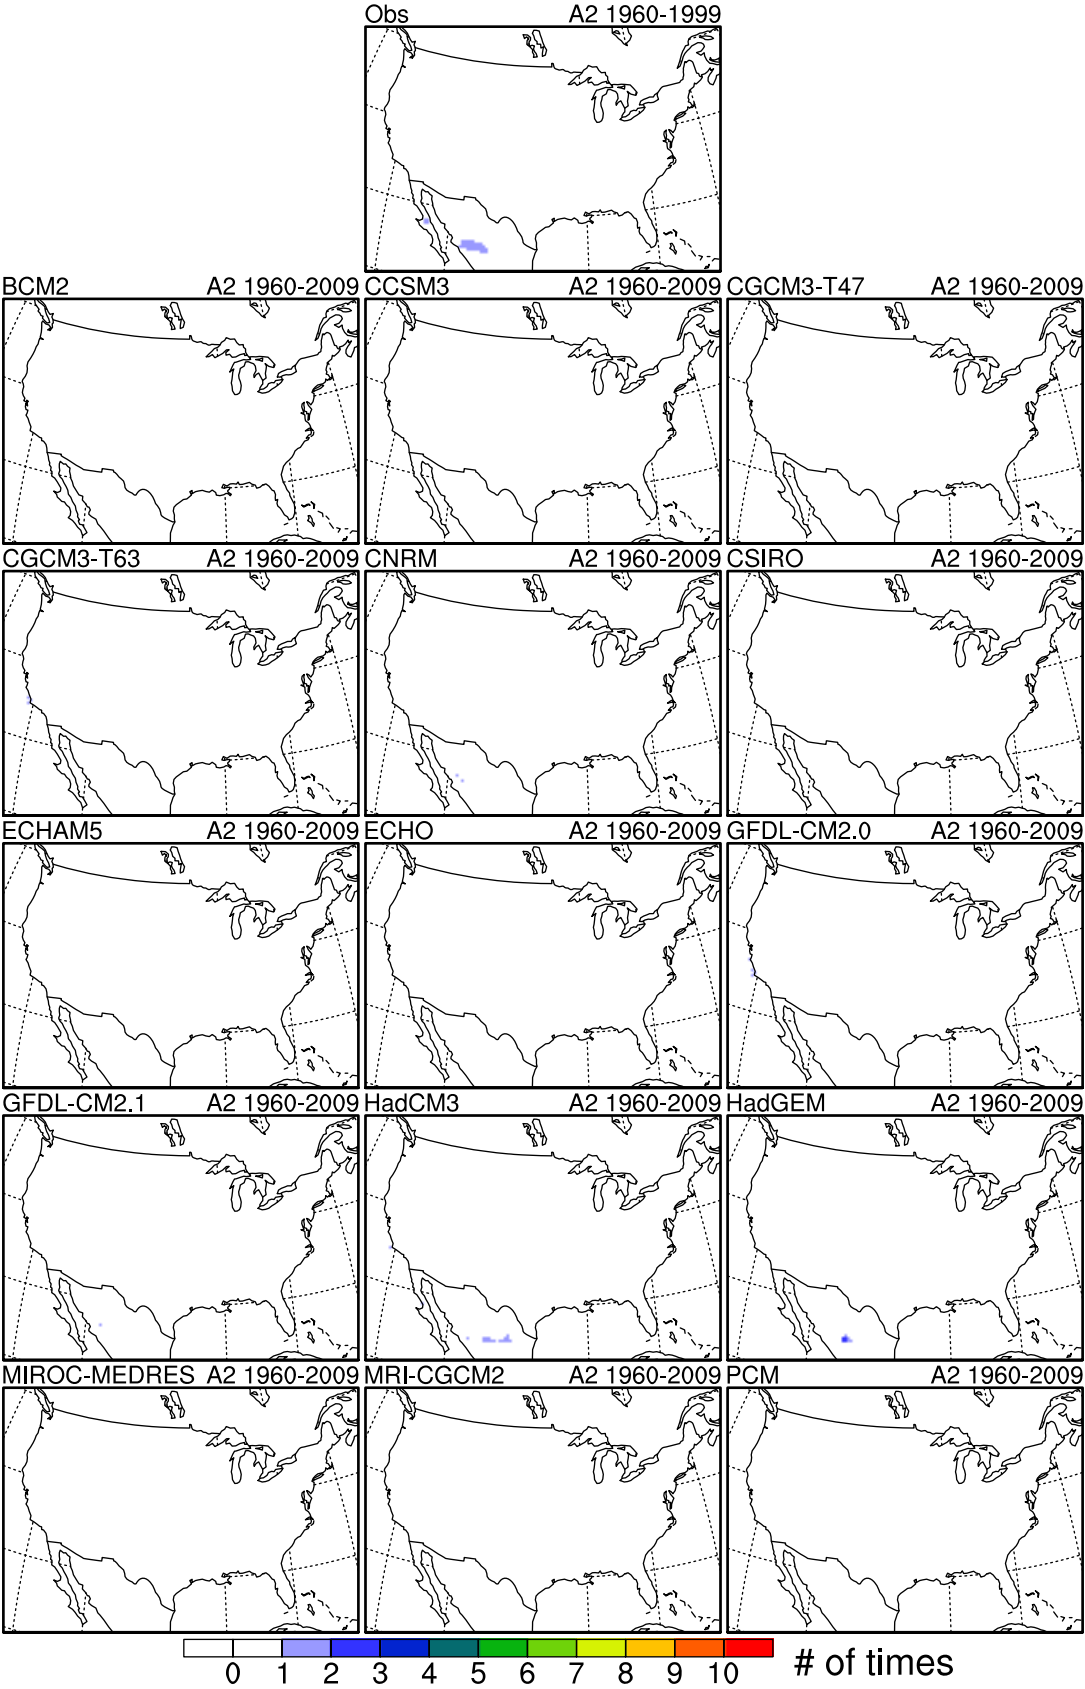
Number of times a tmax value repeats > 3x



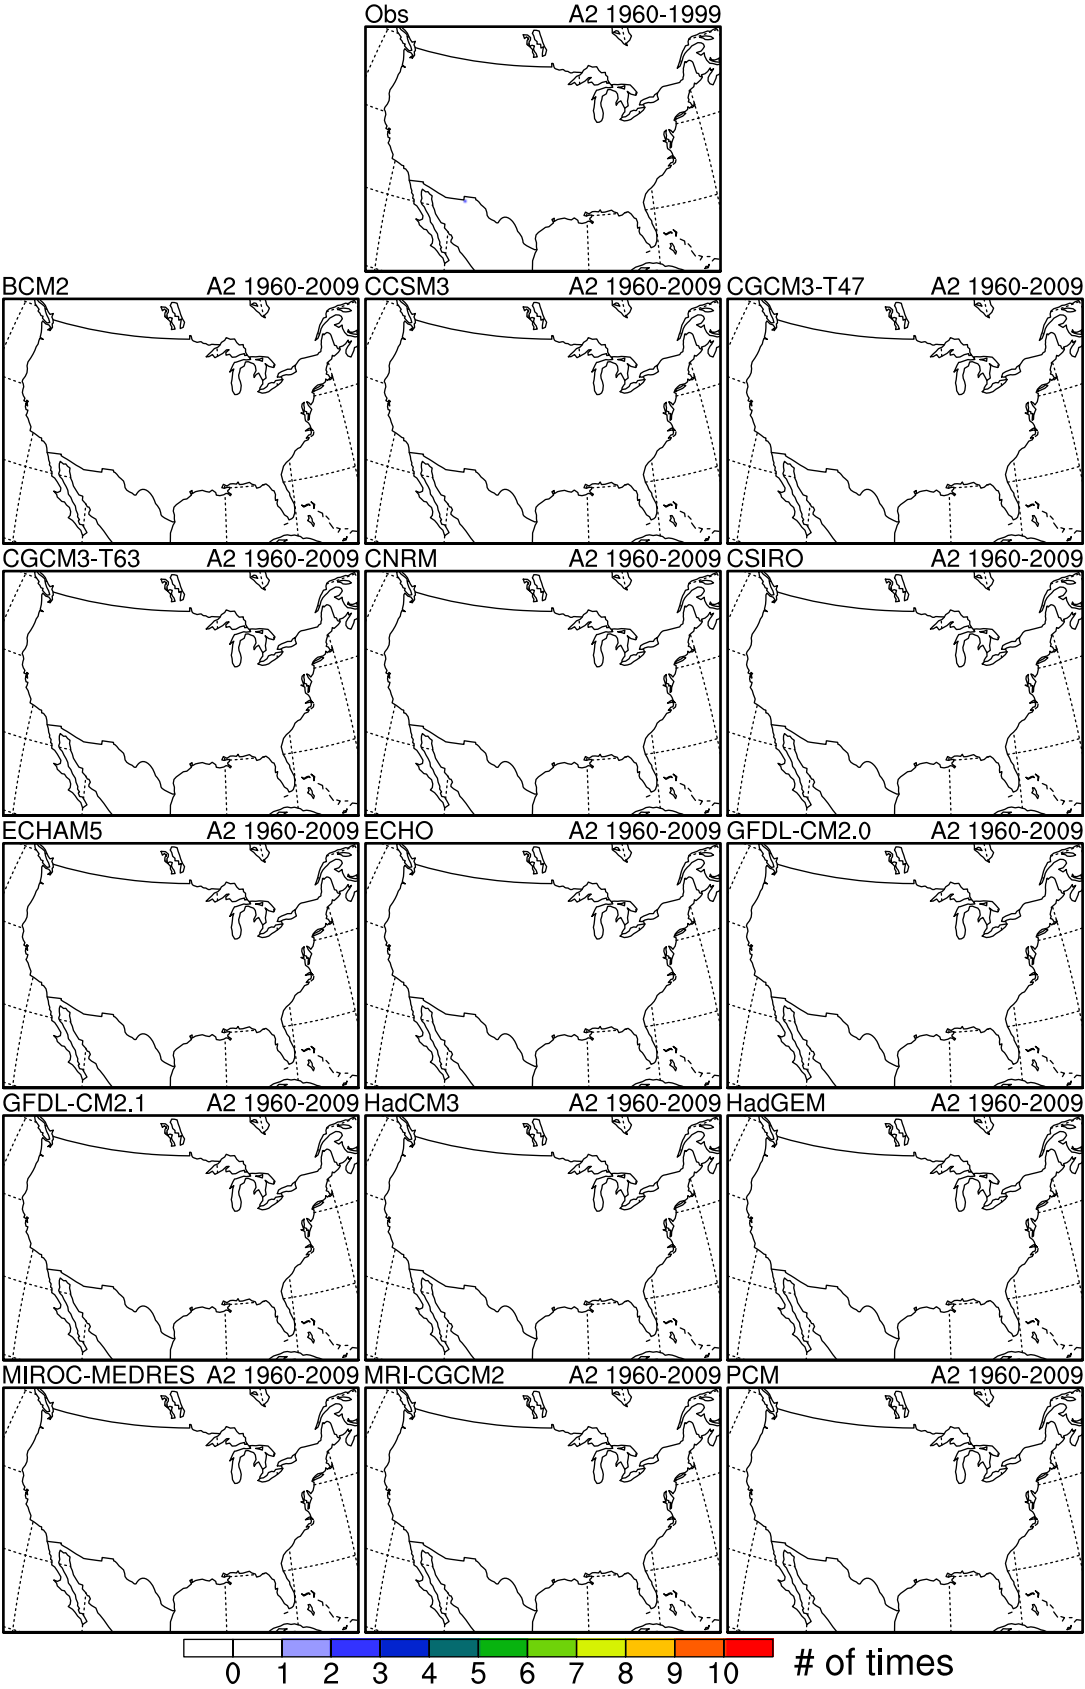
Number of times a tmin value repeats > 3x



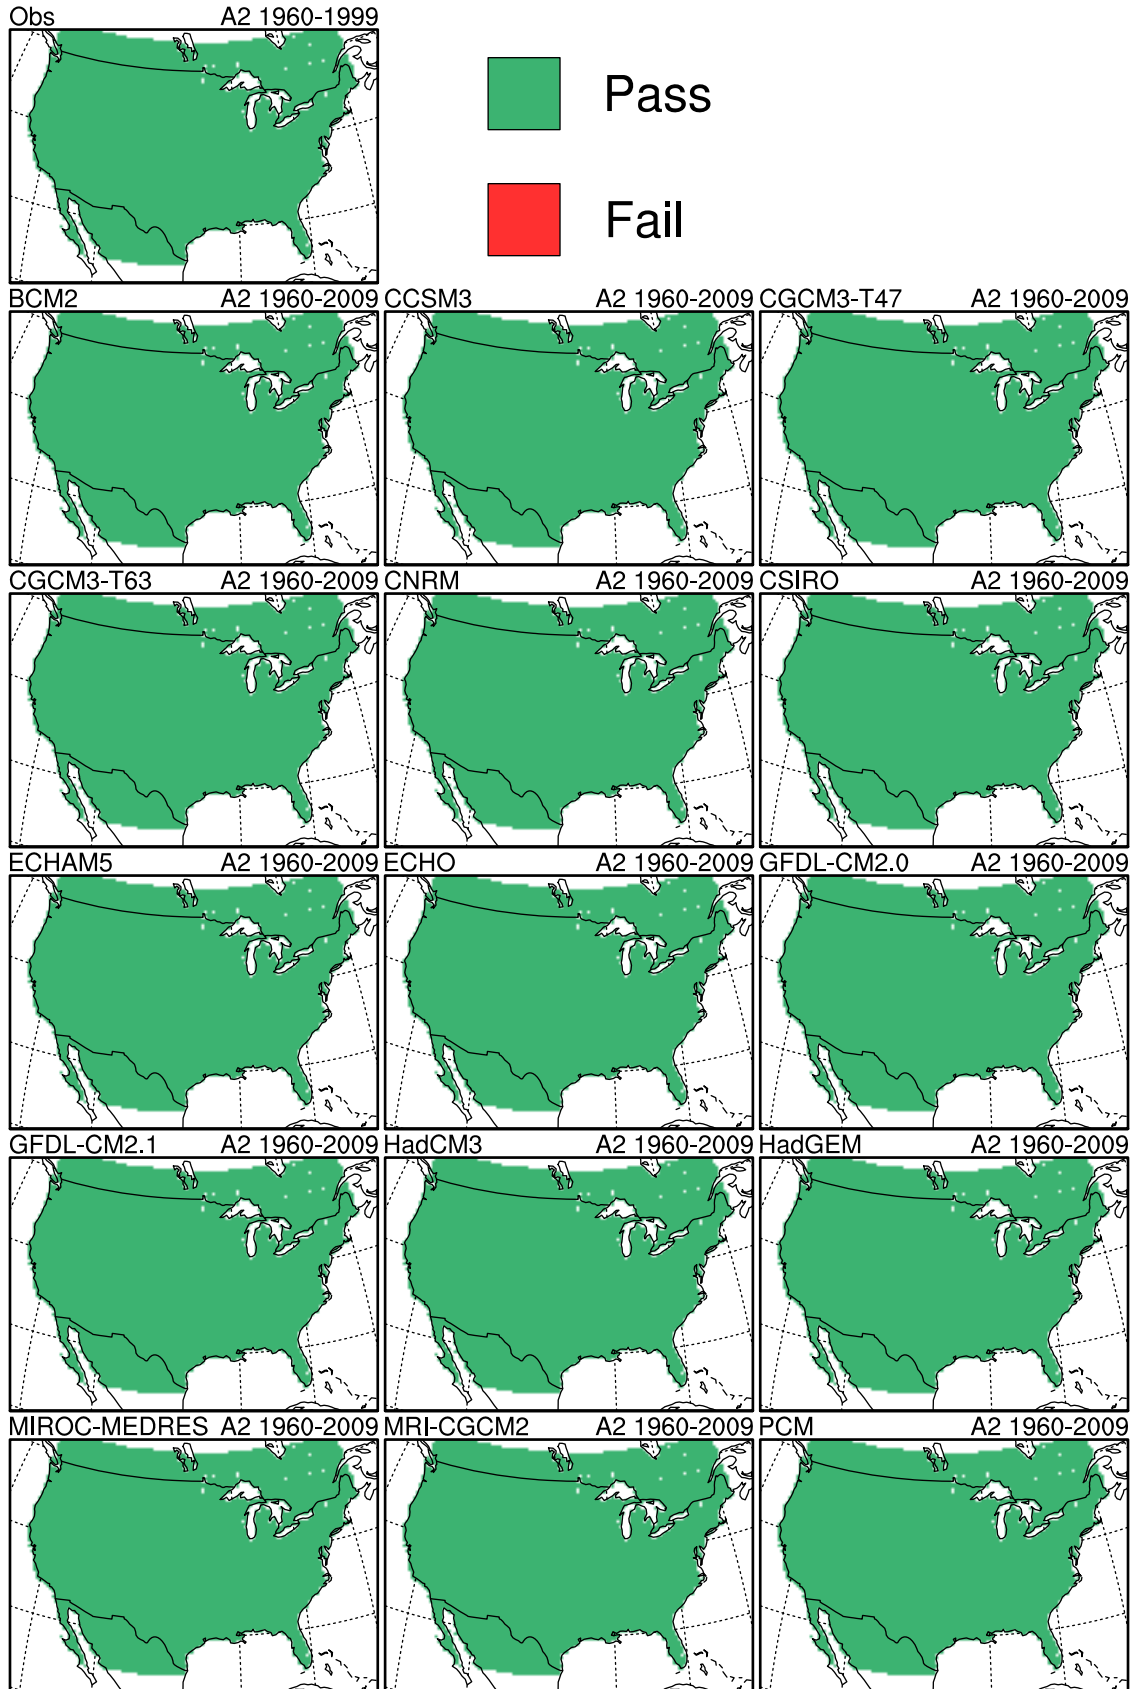
Number of times winter (Jan-Feb) tmax is warmer than summer (Jul-Aug)



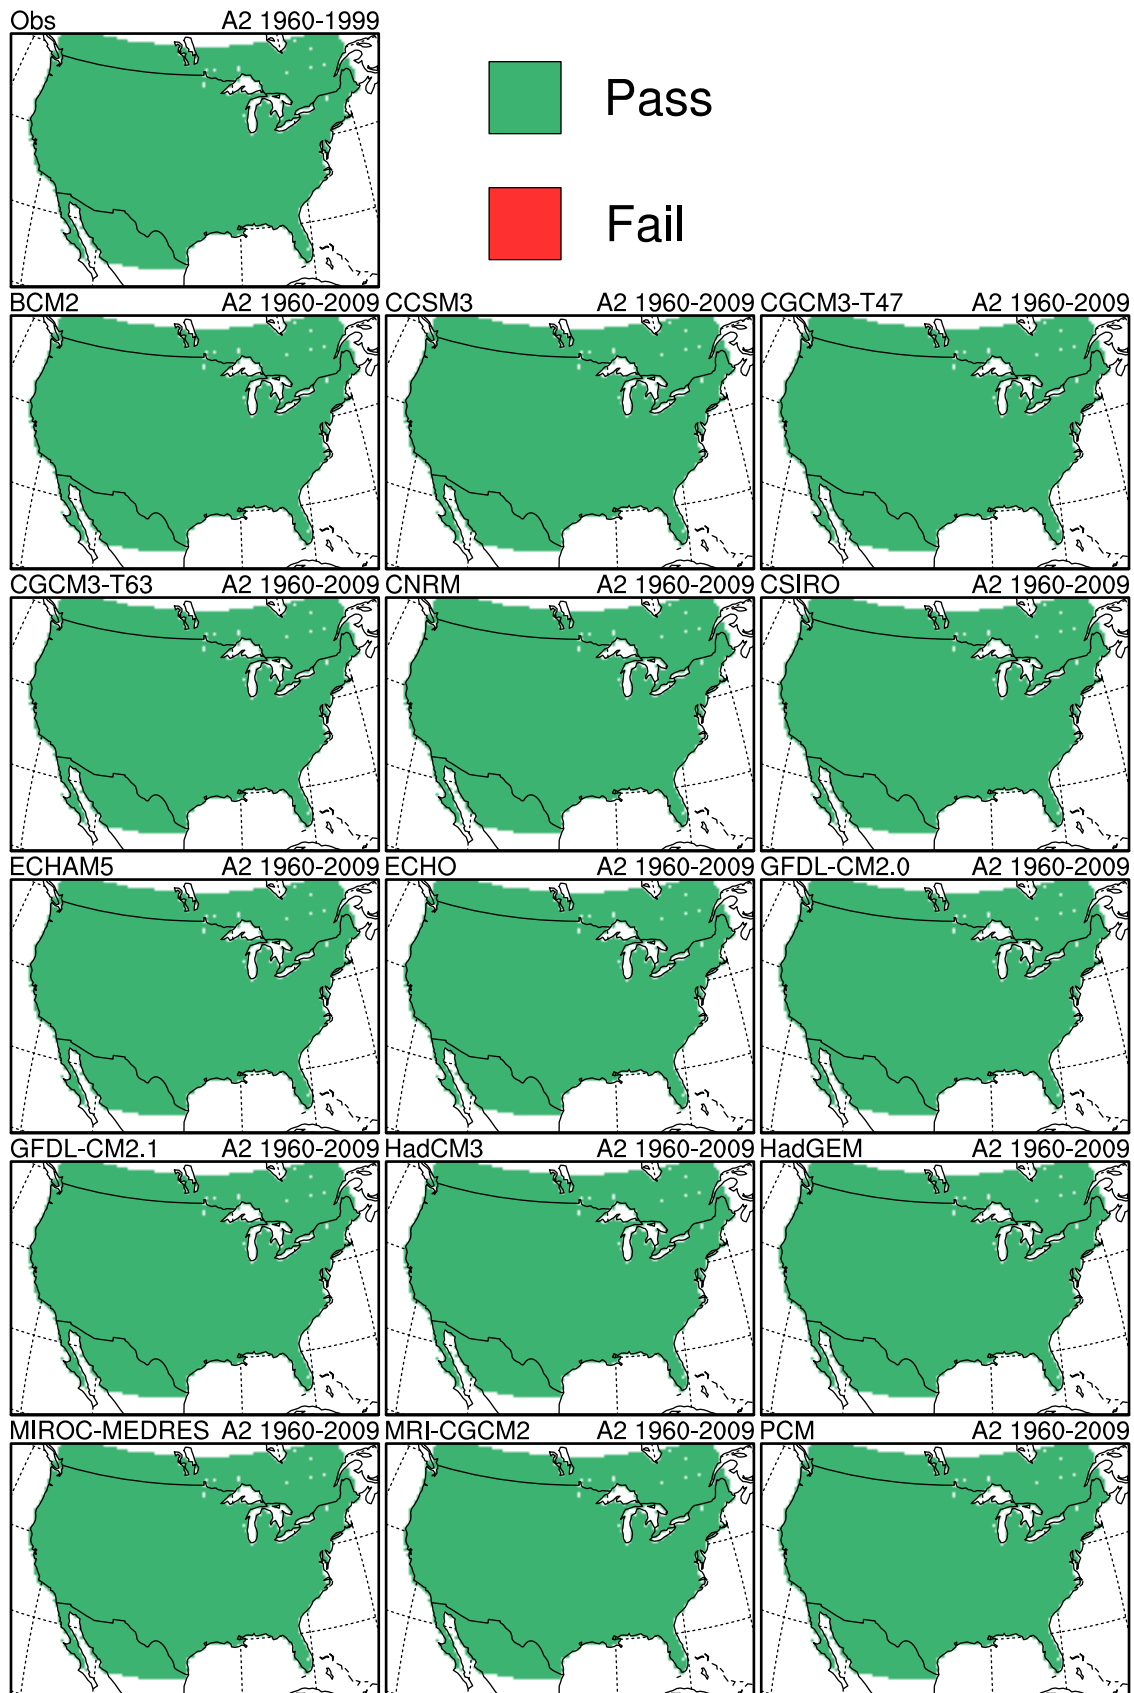
Number of times winter (Jan-Feb) tmin is warmer than summer (Jul-Aug)




Occurrences of $\max(t_{\max}) < \min$ observed CONUS value



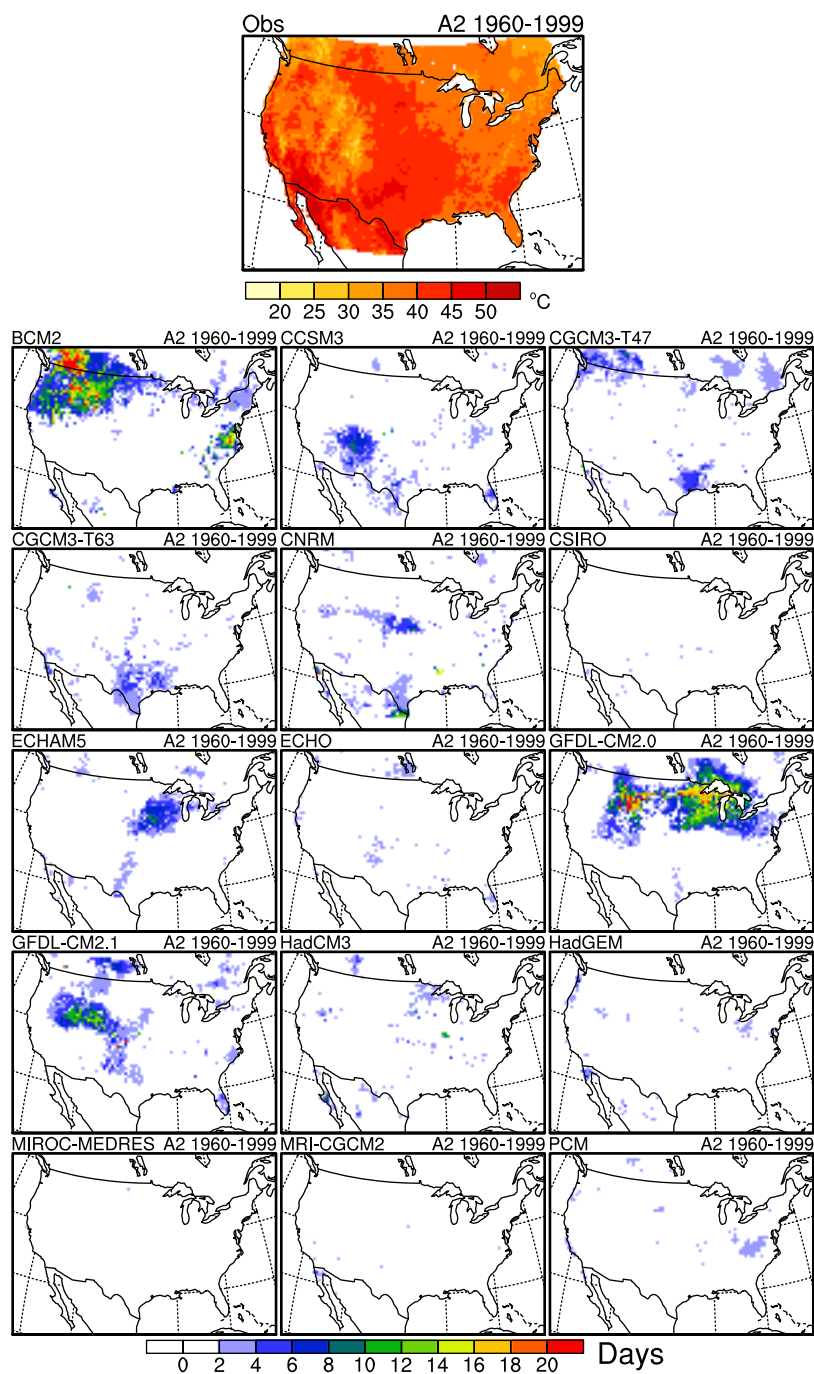
Occurrences of $\min(t_{\min}) > \max \text{ observed CONUS value}$

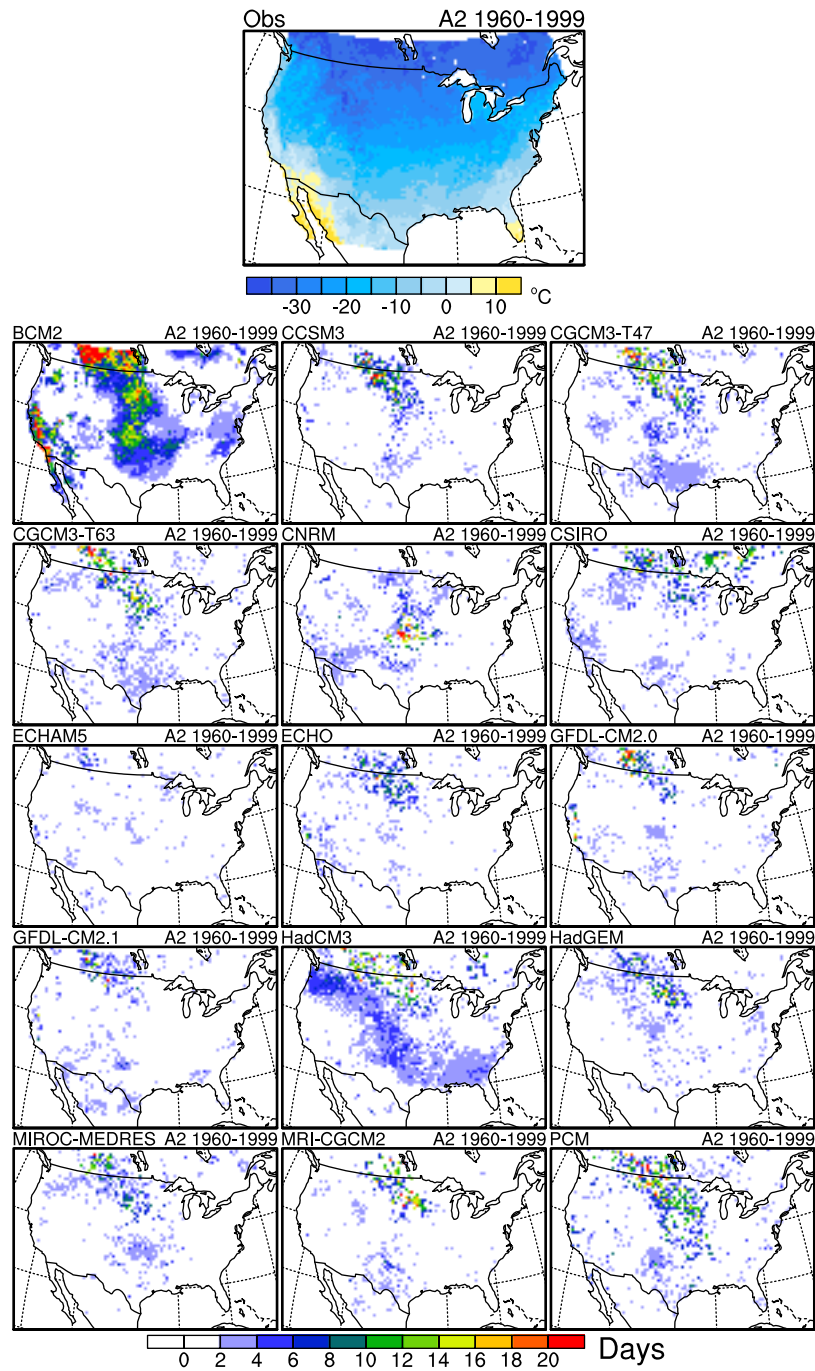


Occurrences of $\min(t_{\min}) < \min$ observed CONUS value

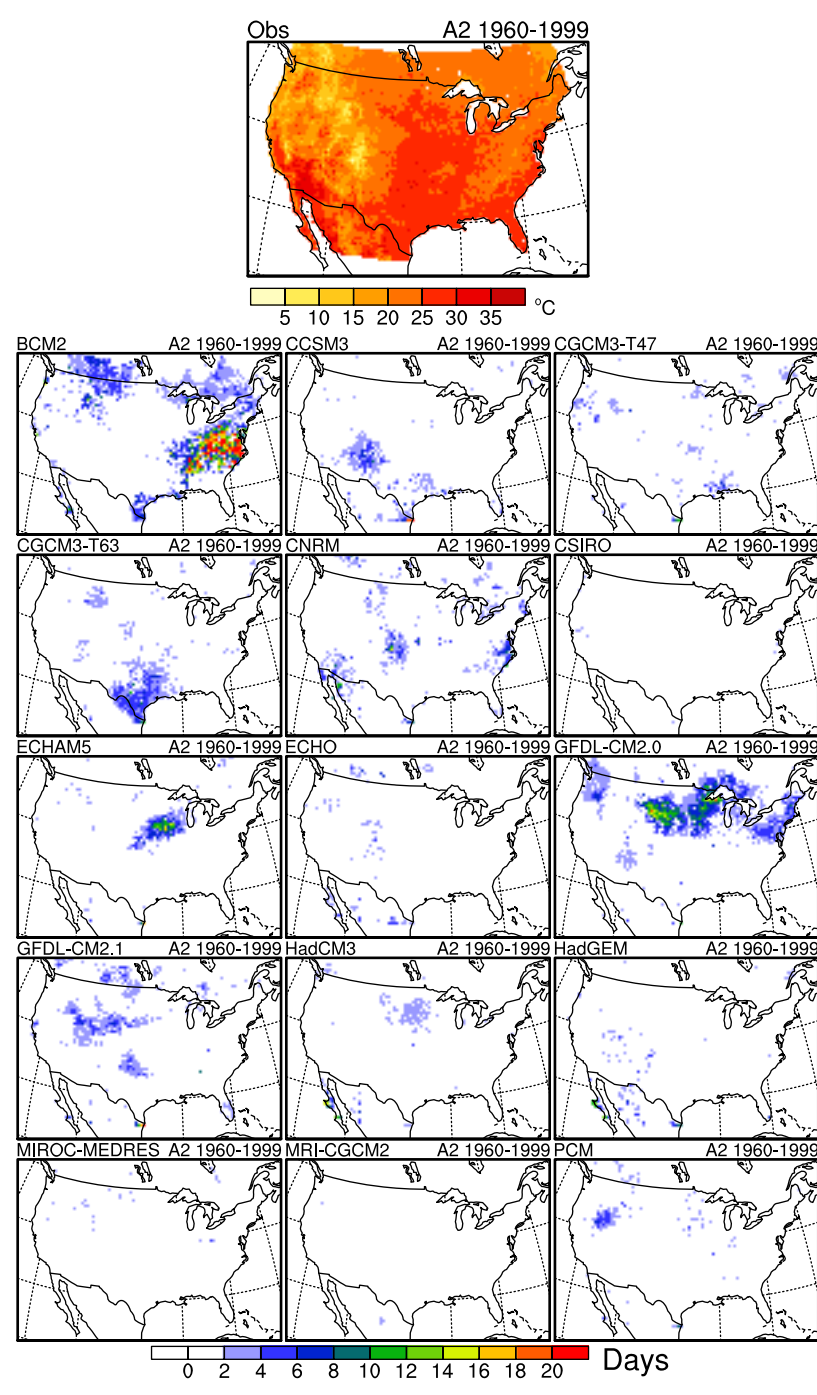


Number of historical days with $t_{max} > \max \text{ observed local } t_{max} \text{ value}$

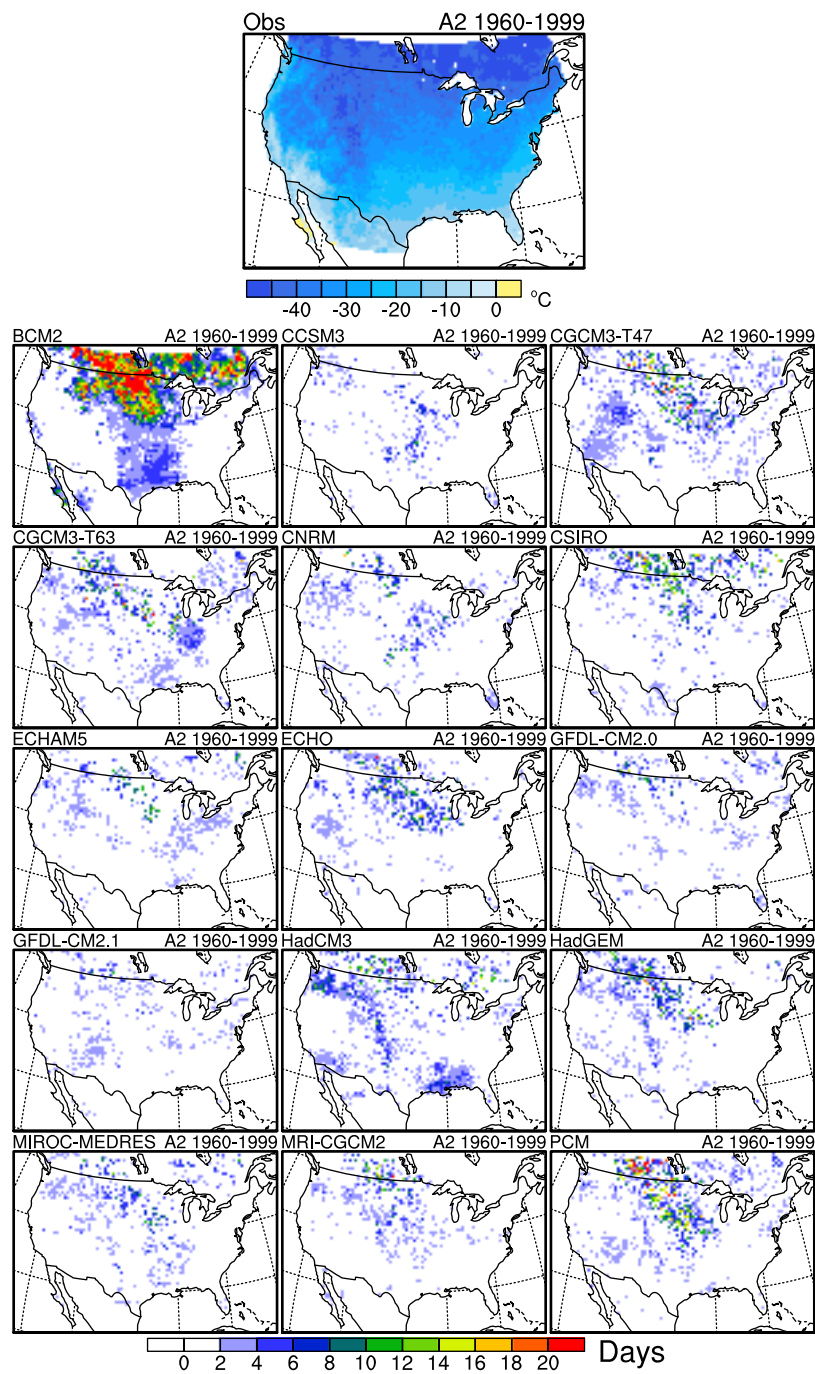


Number of historical days with $t_{max} < \min$ observed local t_{max} value

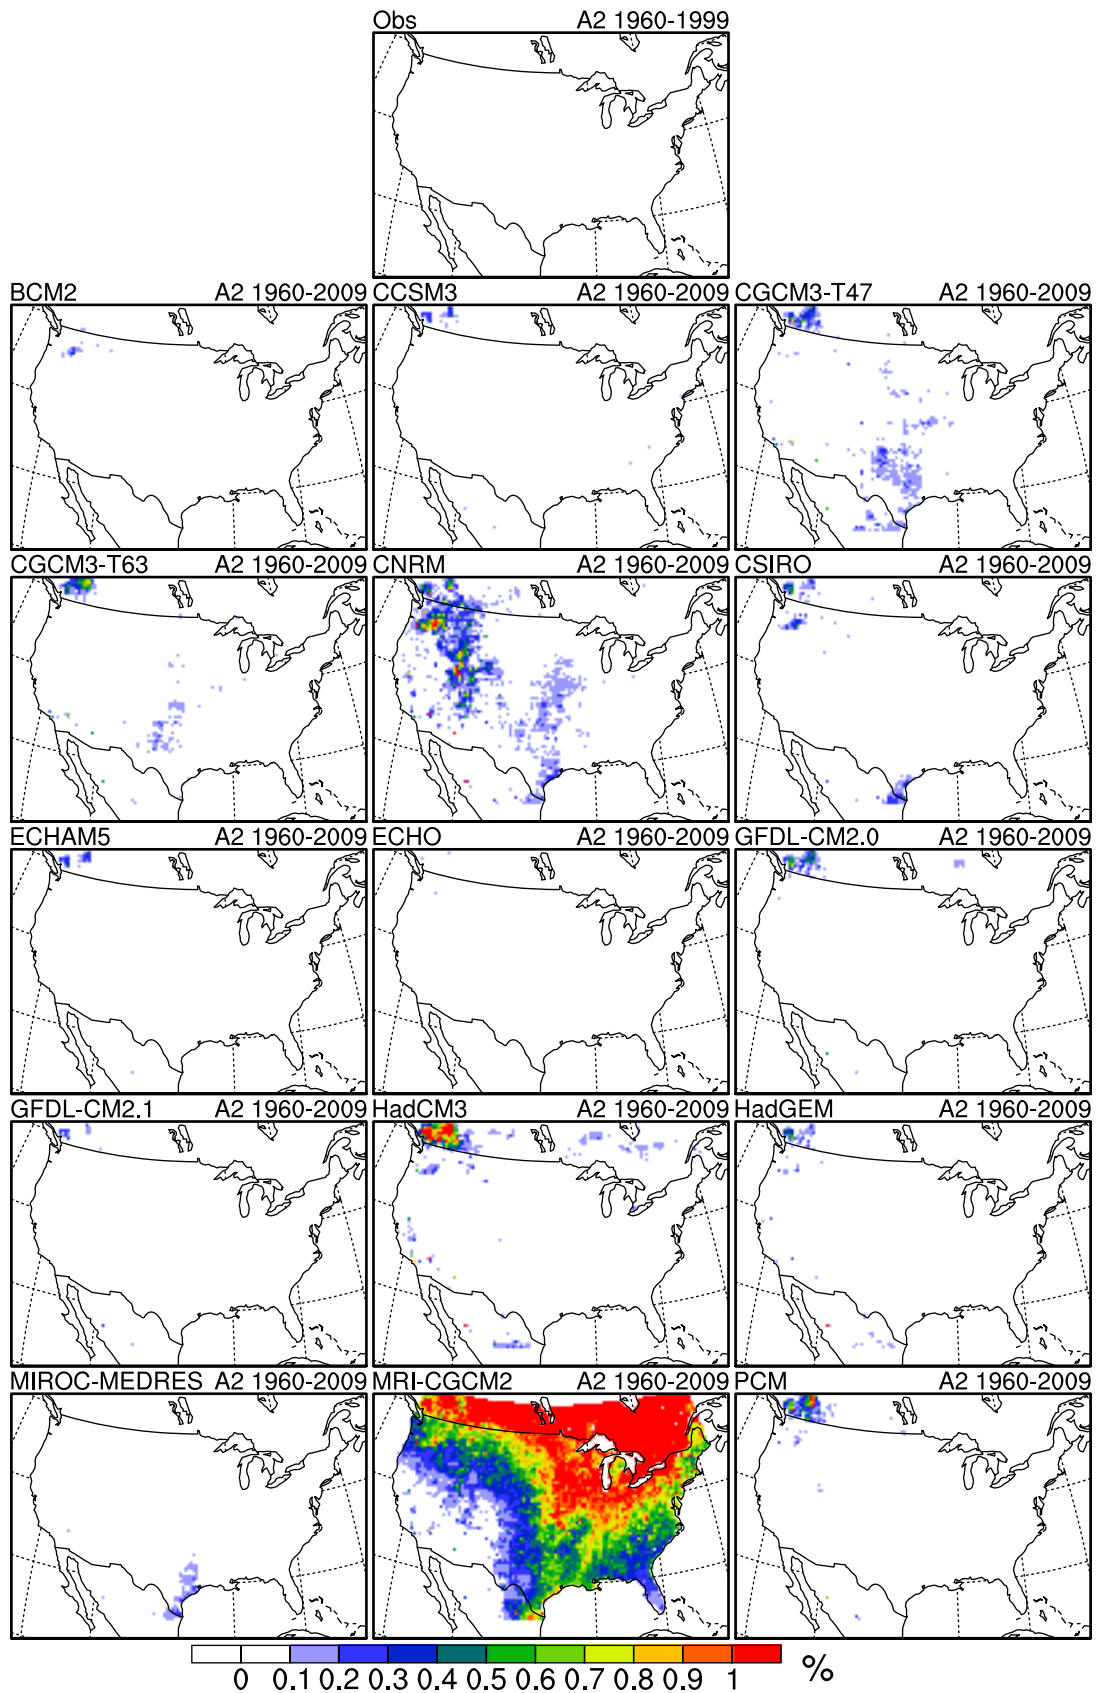
Number of historical days with $t_{min} > \max$ observed local t_{min} value



Number of historical days with $t_{min} < \min$ observed local t_{min} value



Percent (%) of total days where $t_{min} > t_{max}$

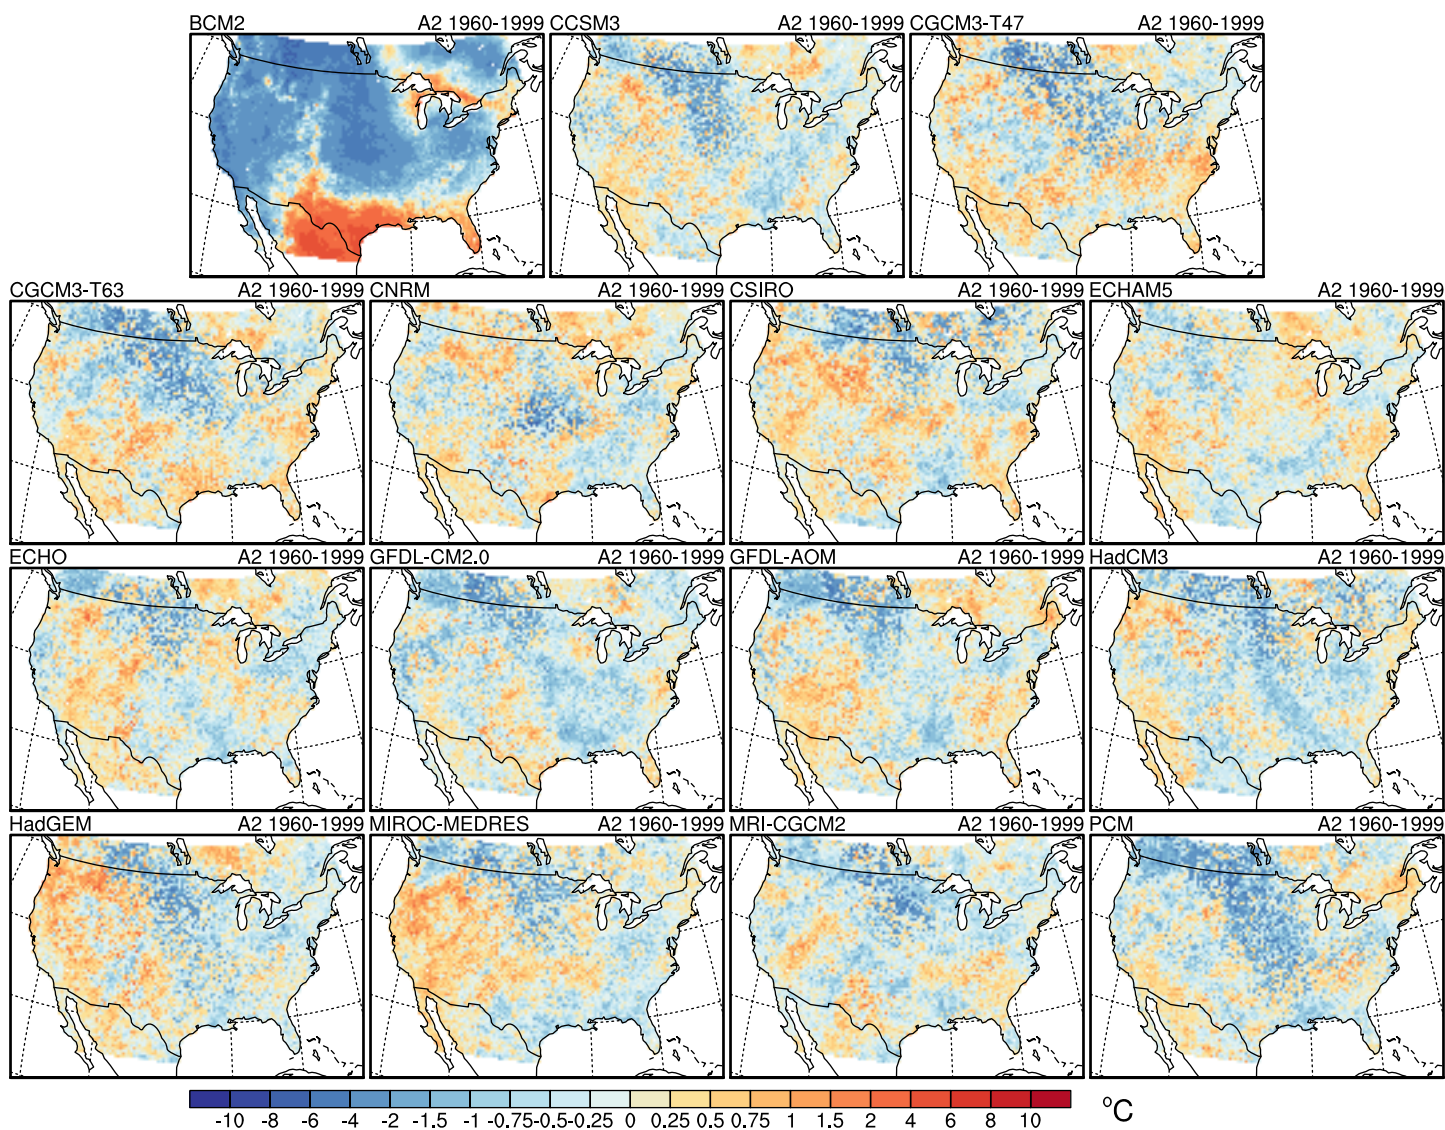


APPENDIX E. Bias Plots

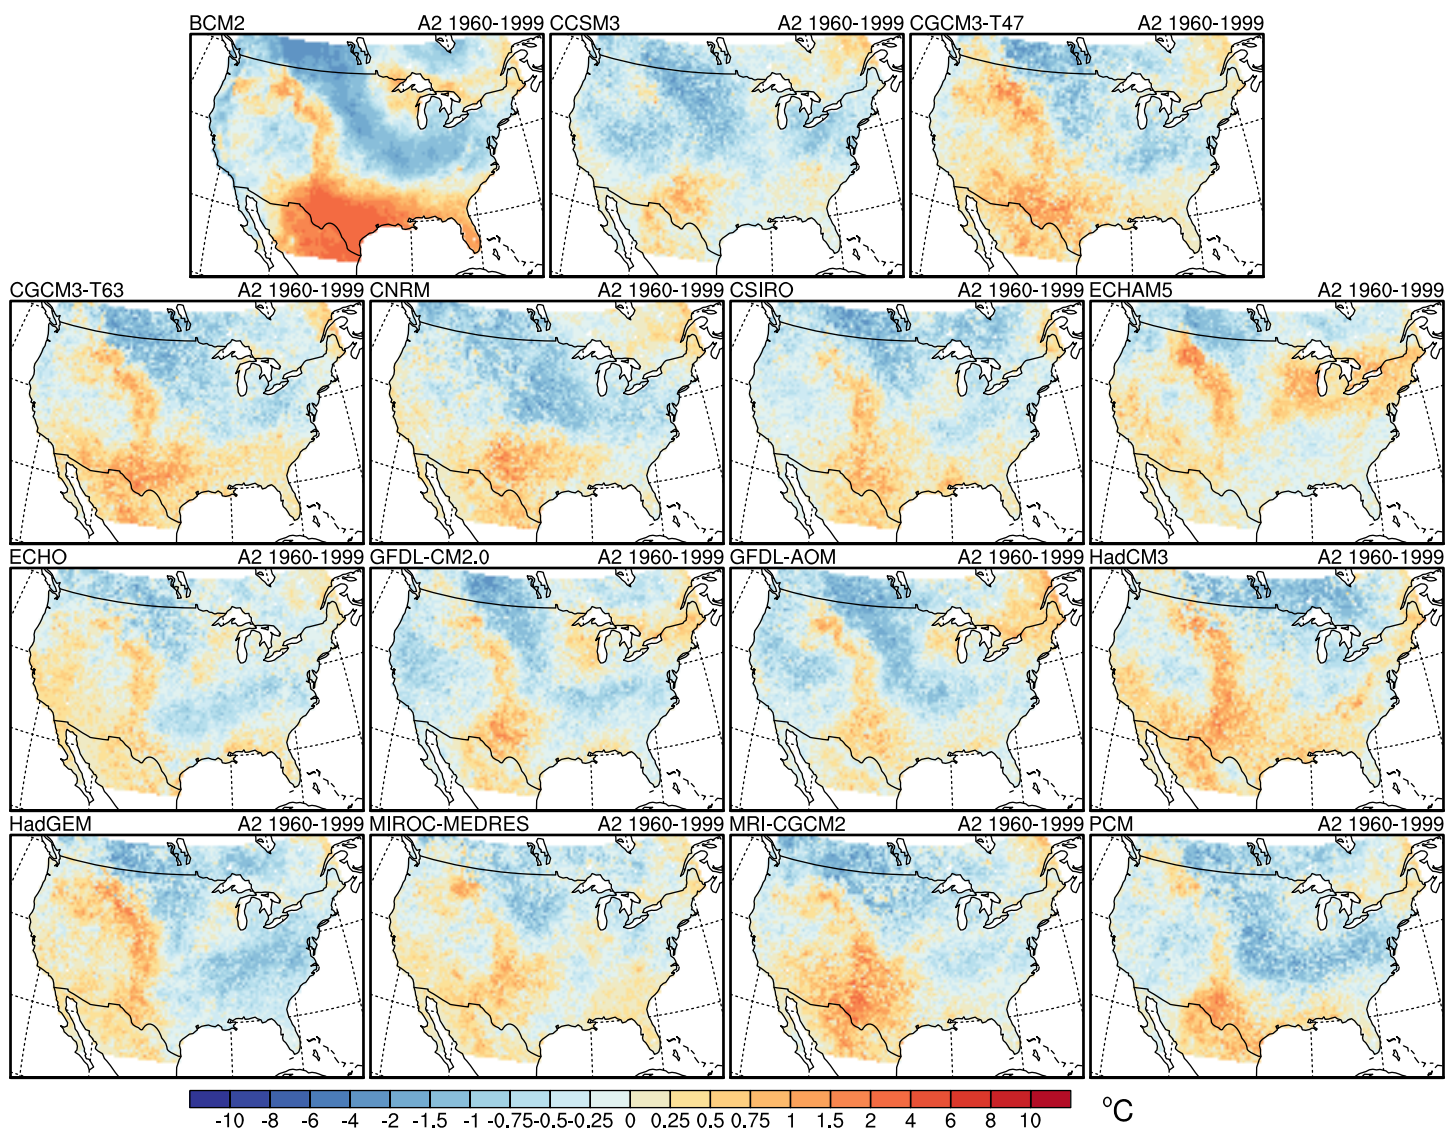
This Appendix provides maps of temperature and precipitation bias tests in the historical gridded downscaling simulations compared to observations over the same time period (1960-1999). The bias tests are described in Table 6.

Original postscript files are also provided in electronic format.

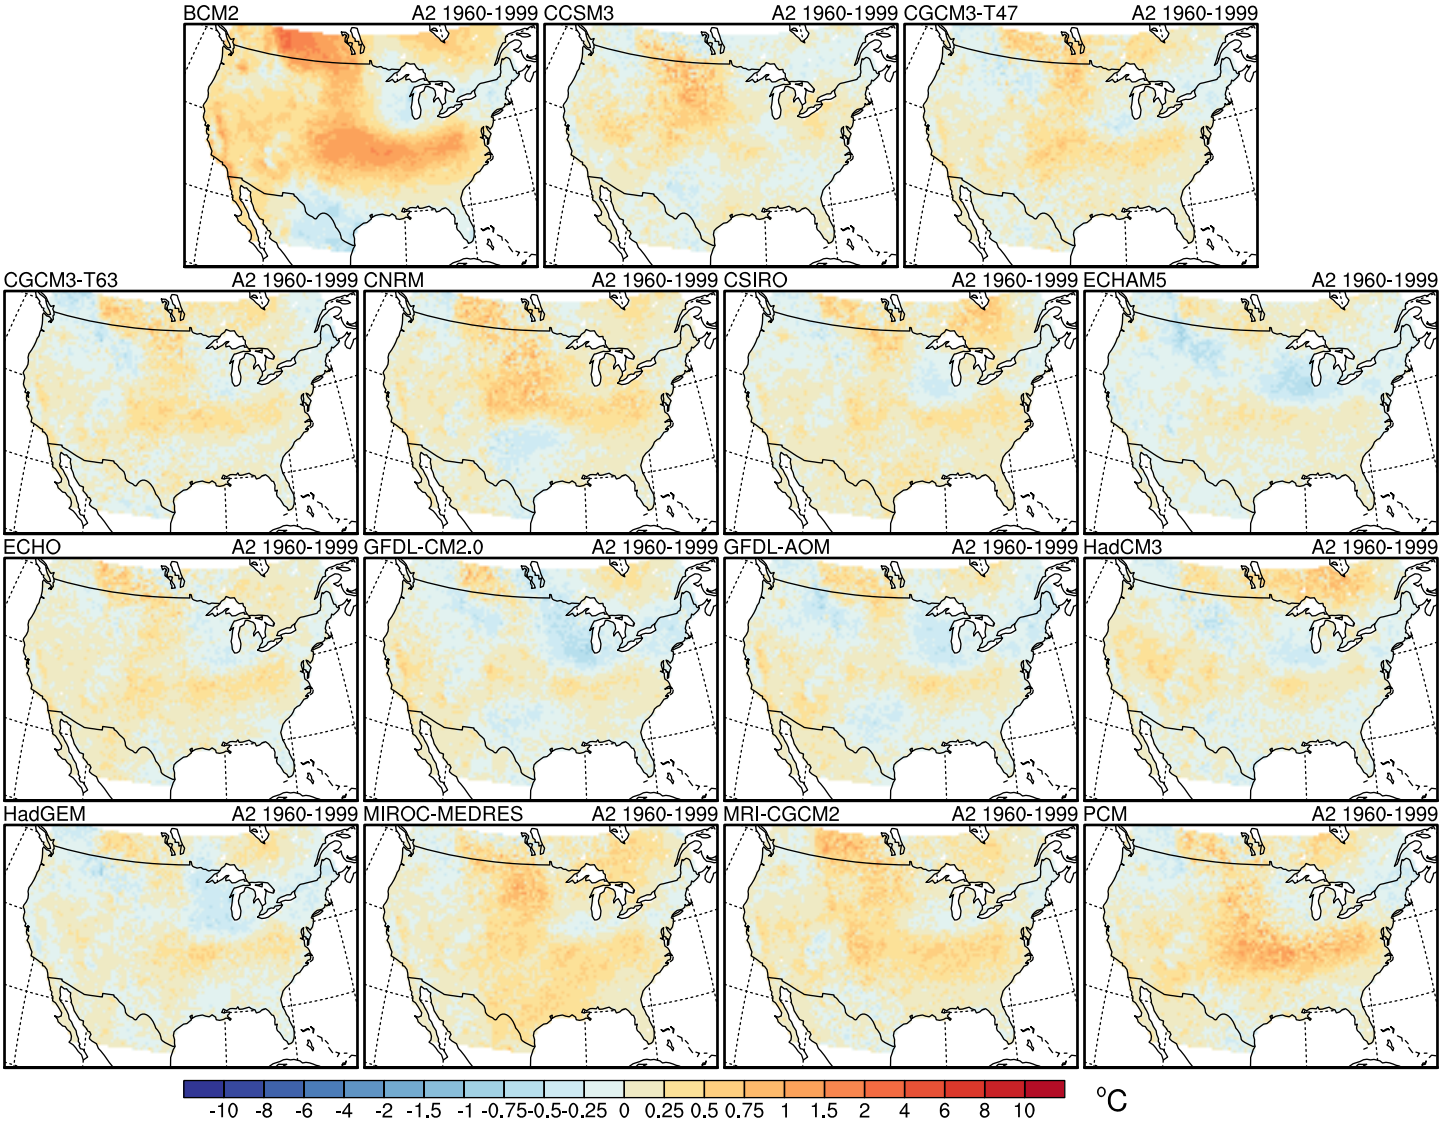
Tmax Bias in 0.1th Quantile



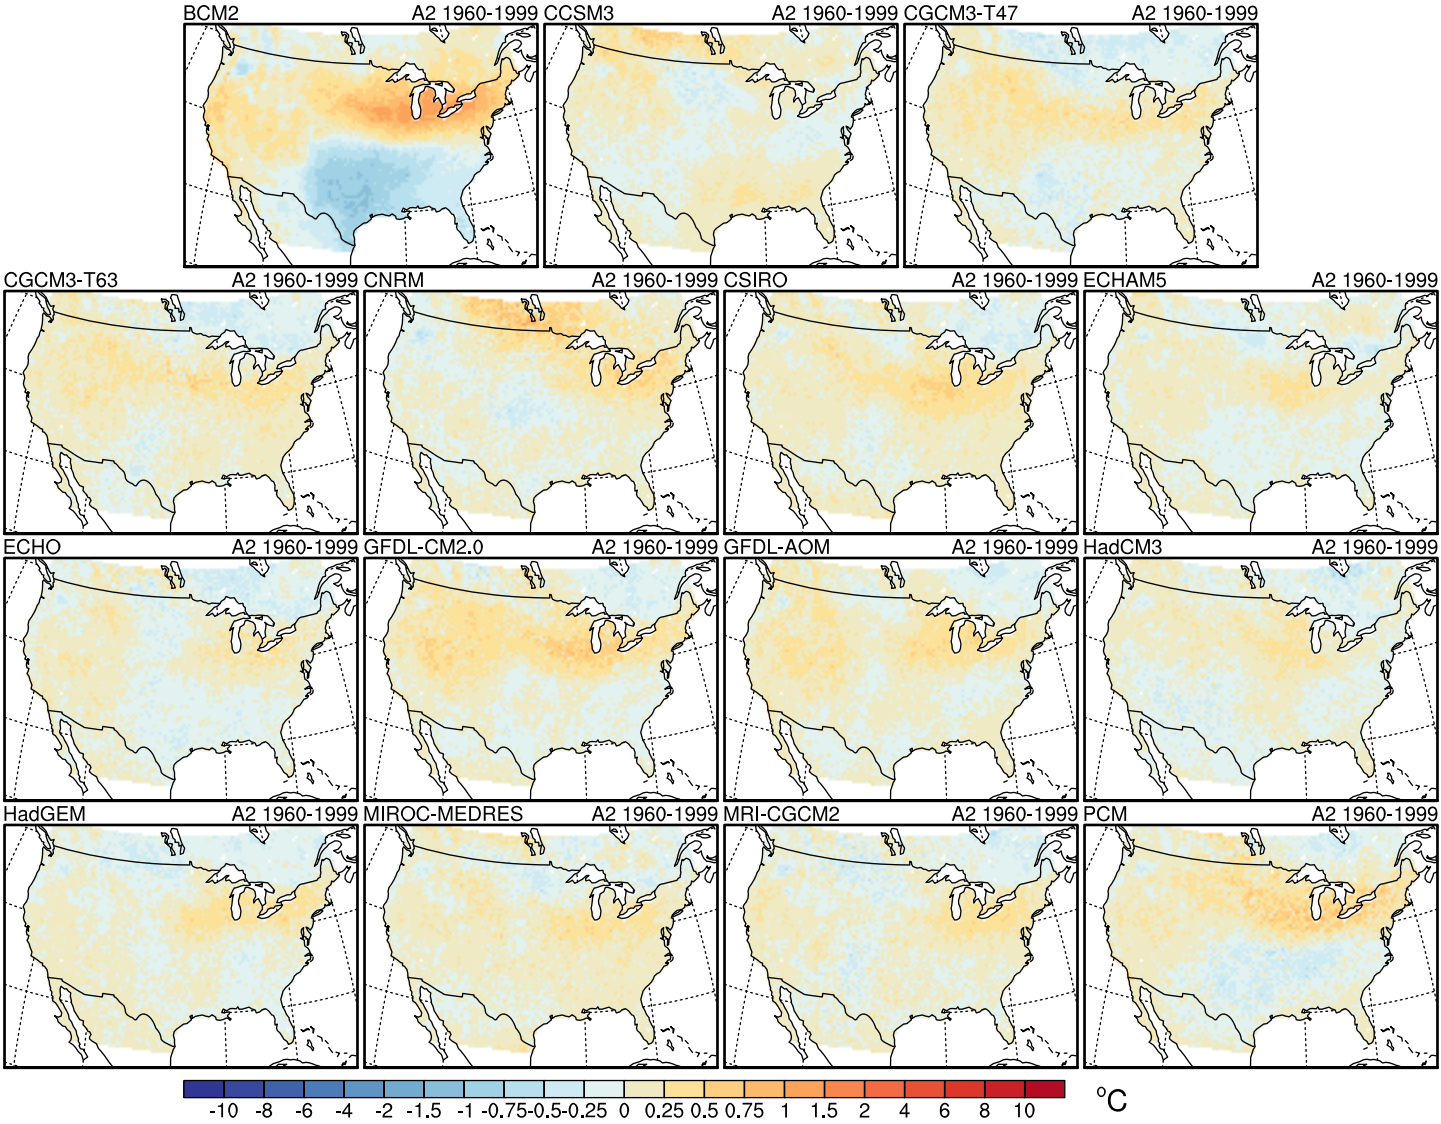
Tmax Bias in 1st Quantile



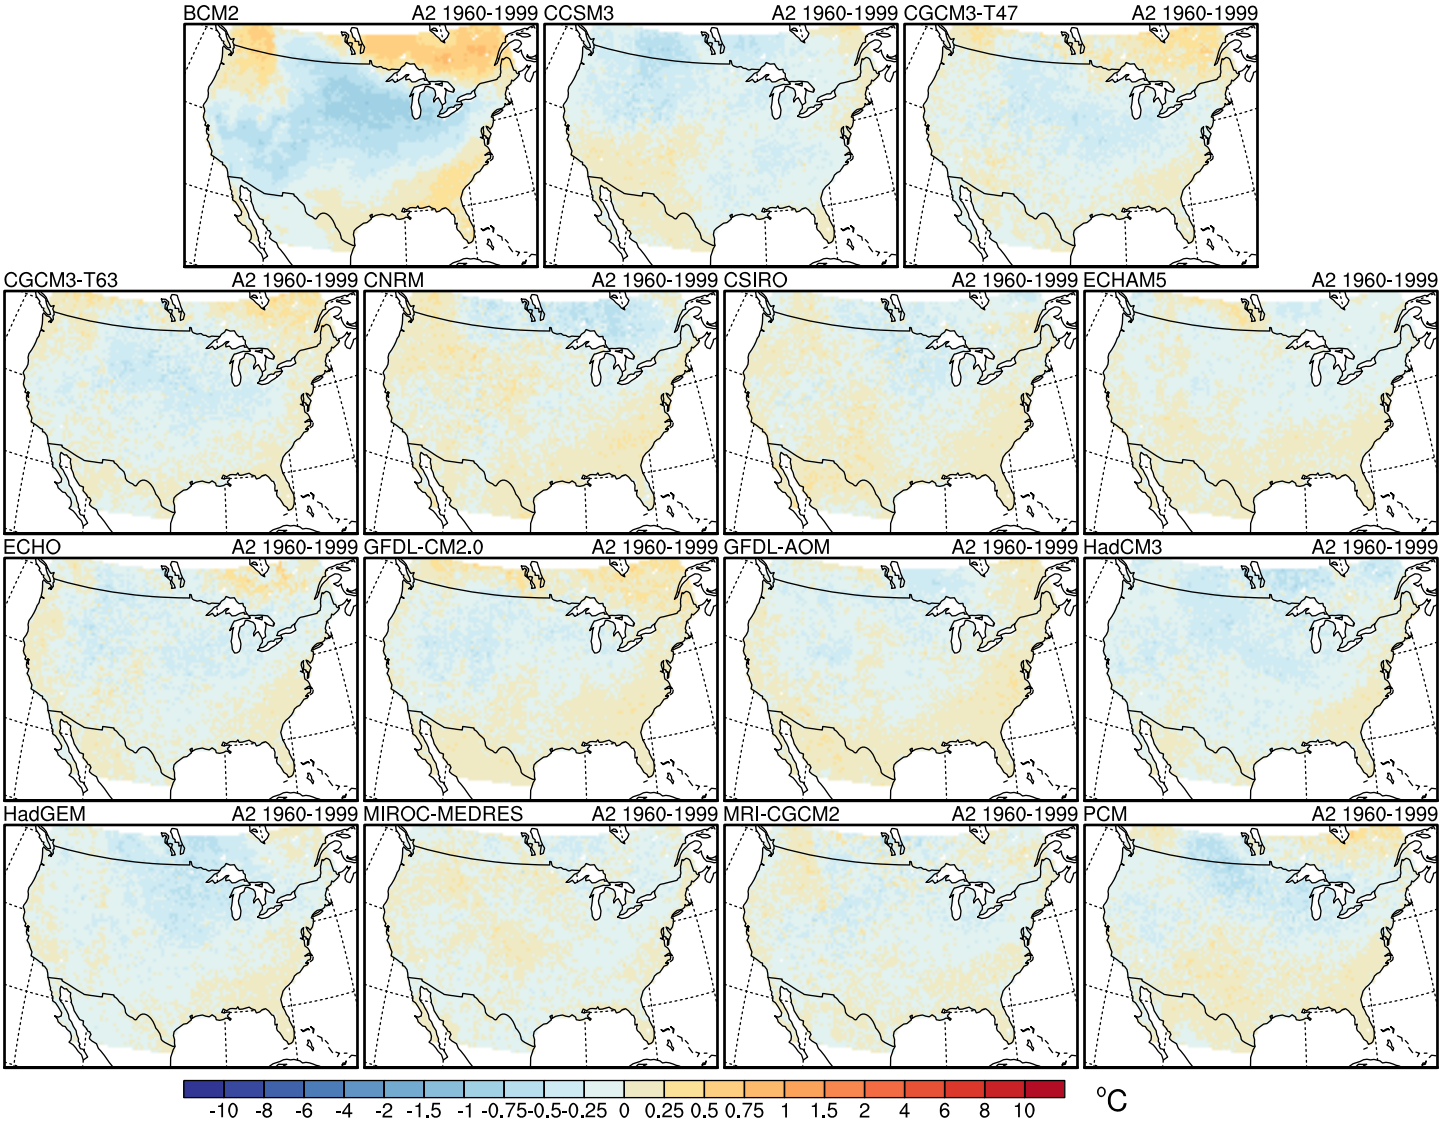
Tmax Bias in 10th Quantile



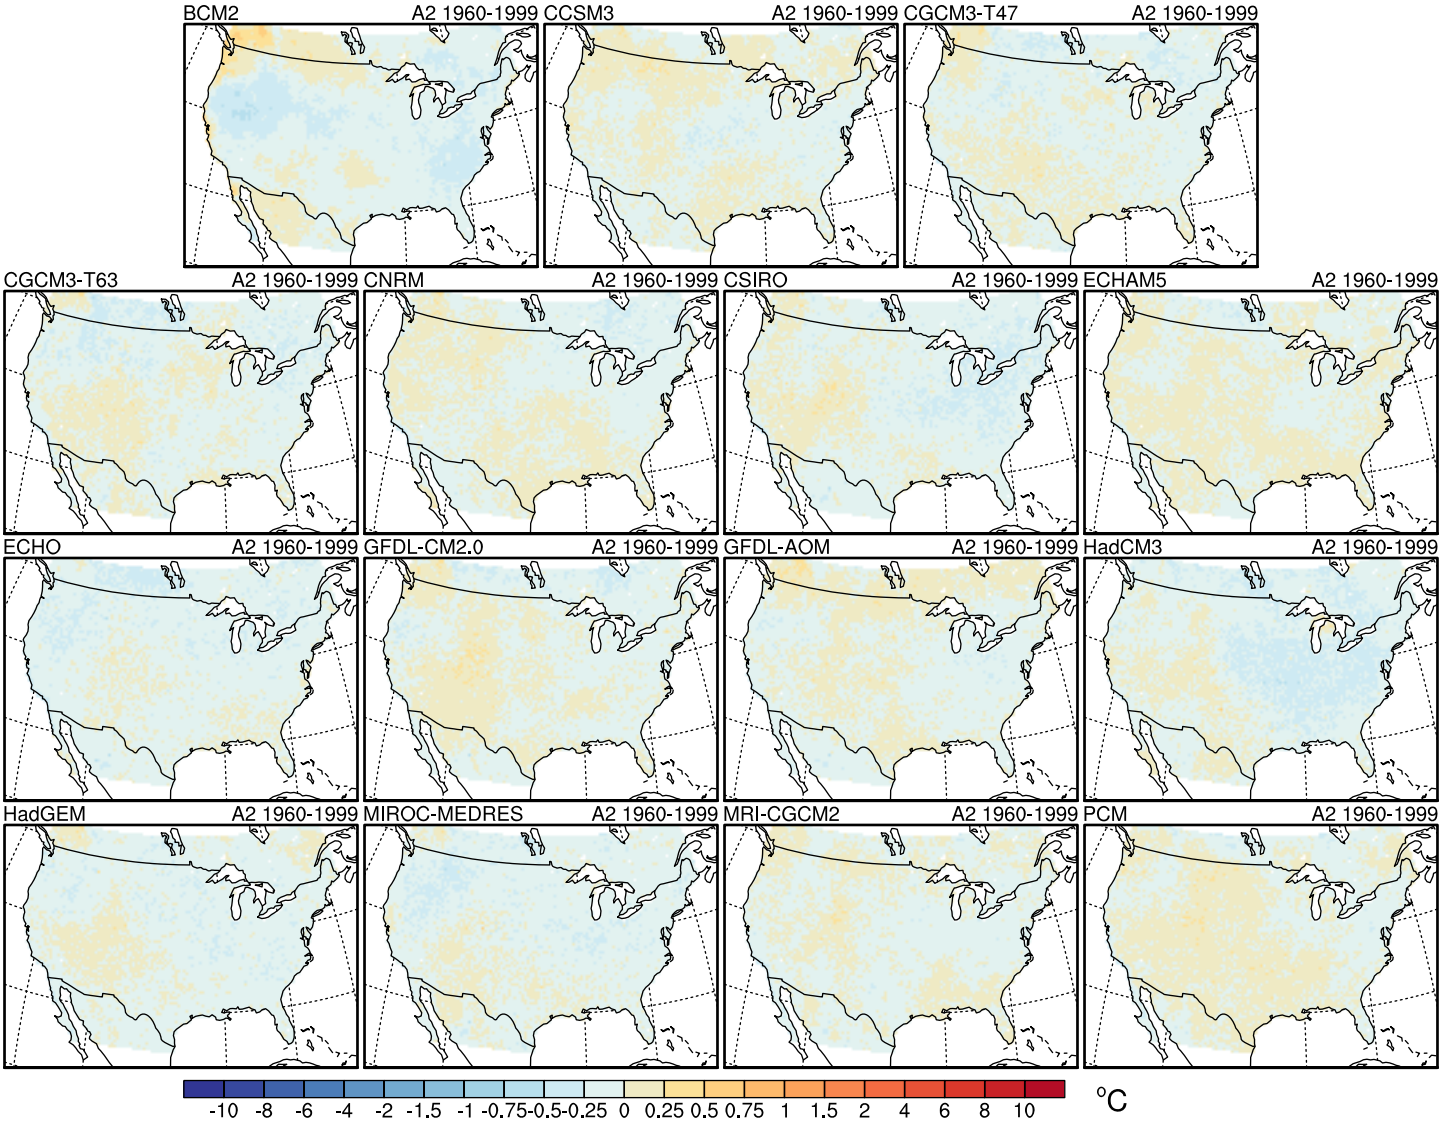
Tmax Bias in 25th Quantile



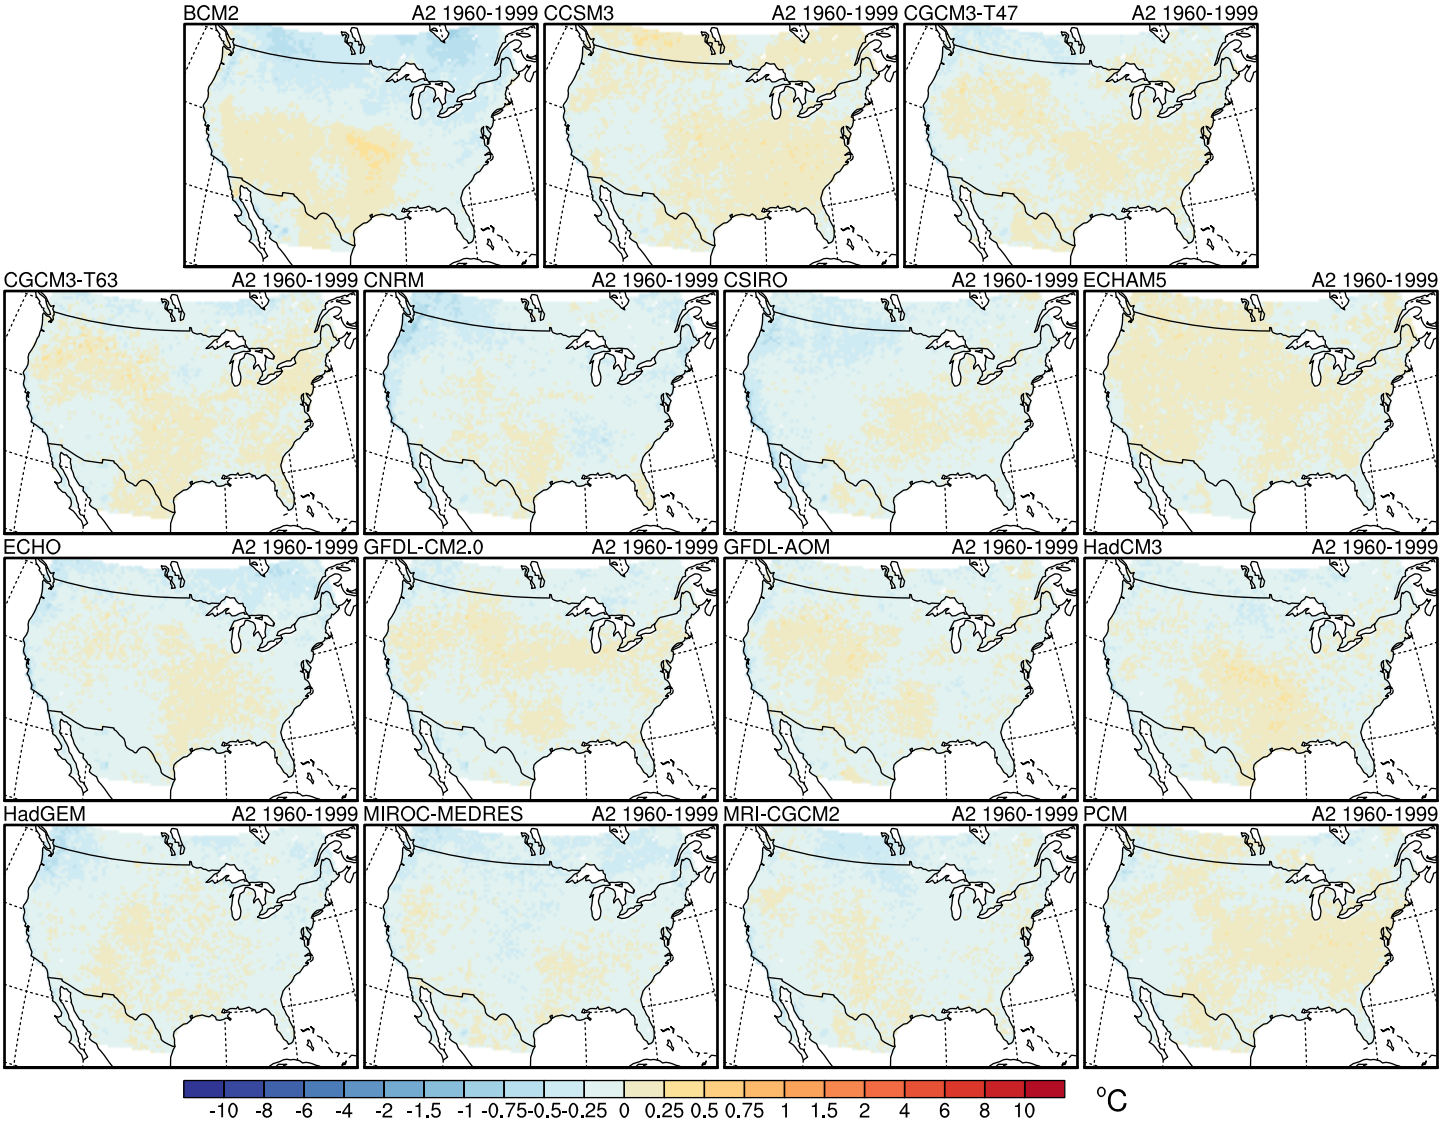
Tmax Bias in 50th Quantile



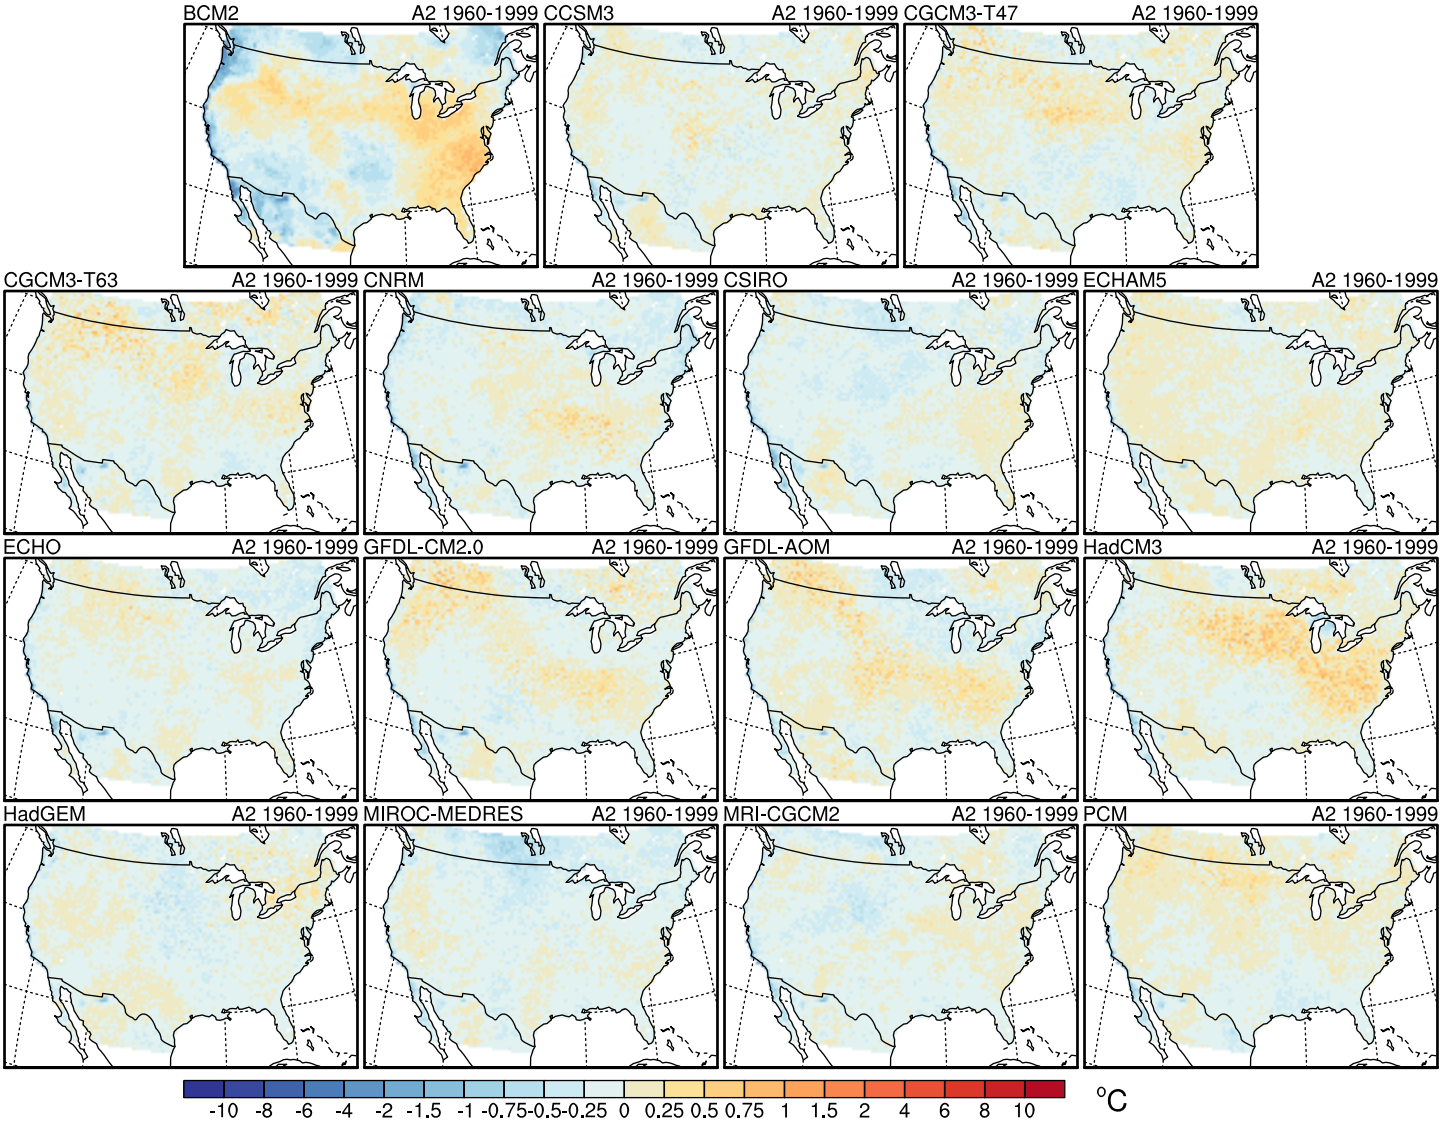
Tmax Bias in 75th Quantile



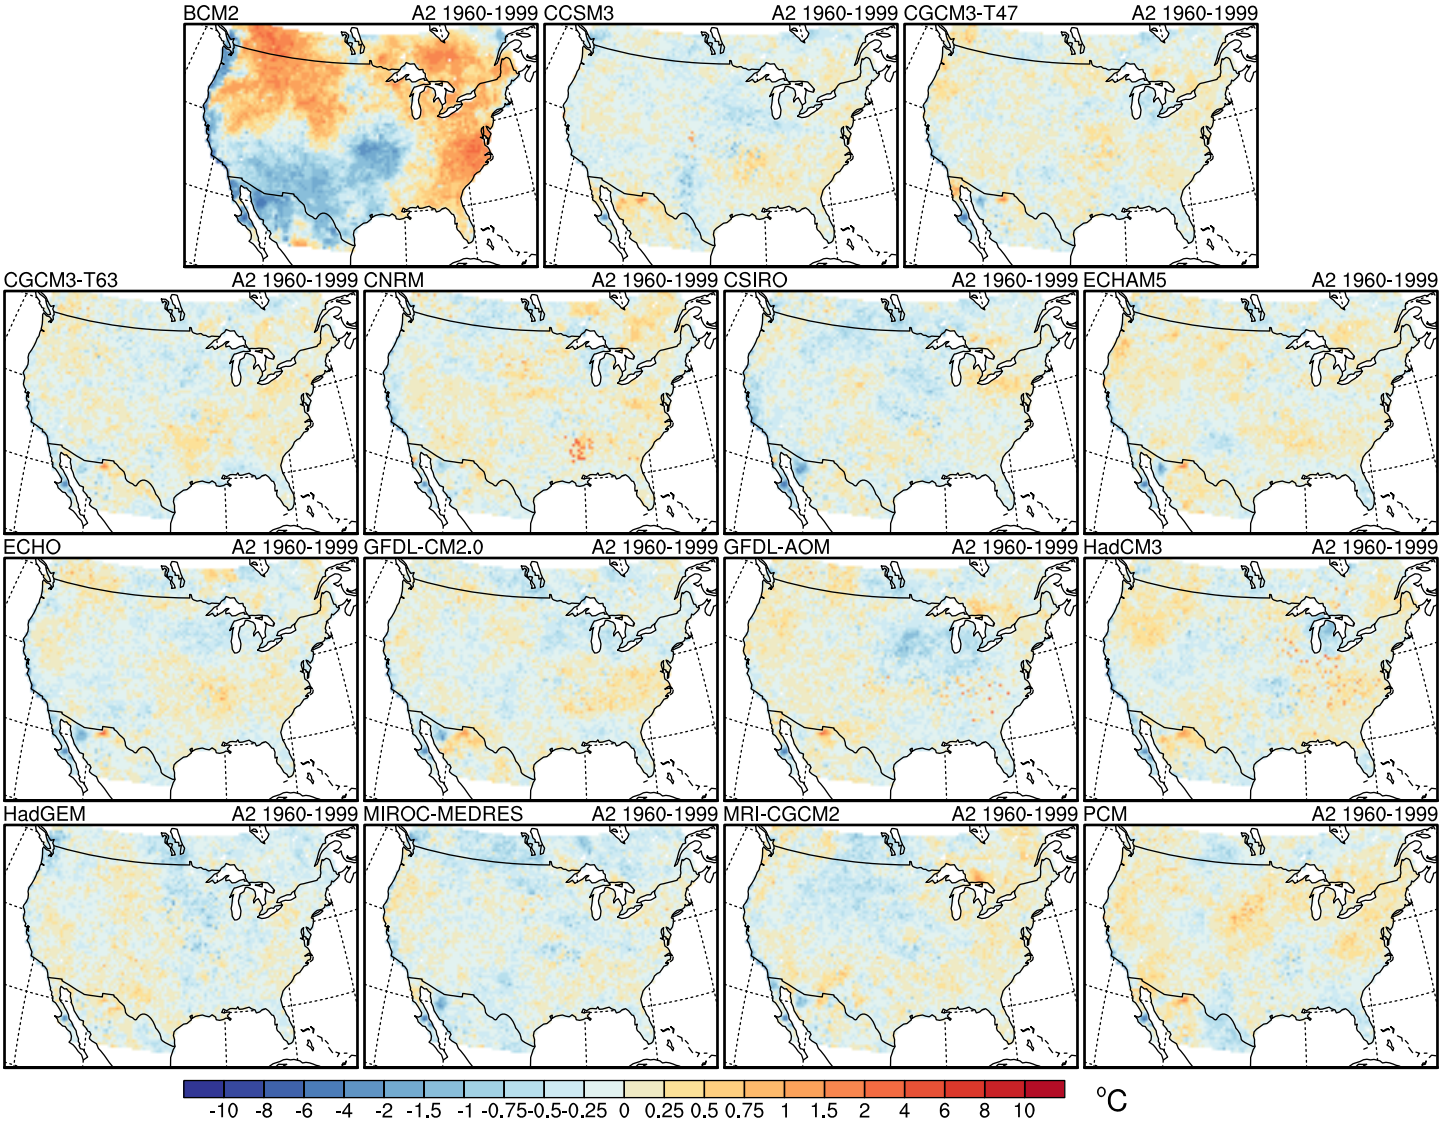
Tmax Bias in 90th Quantile



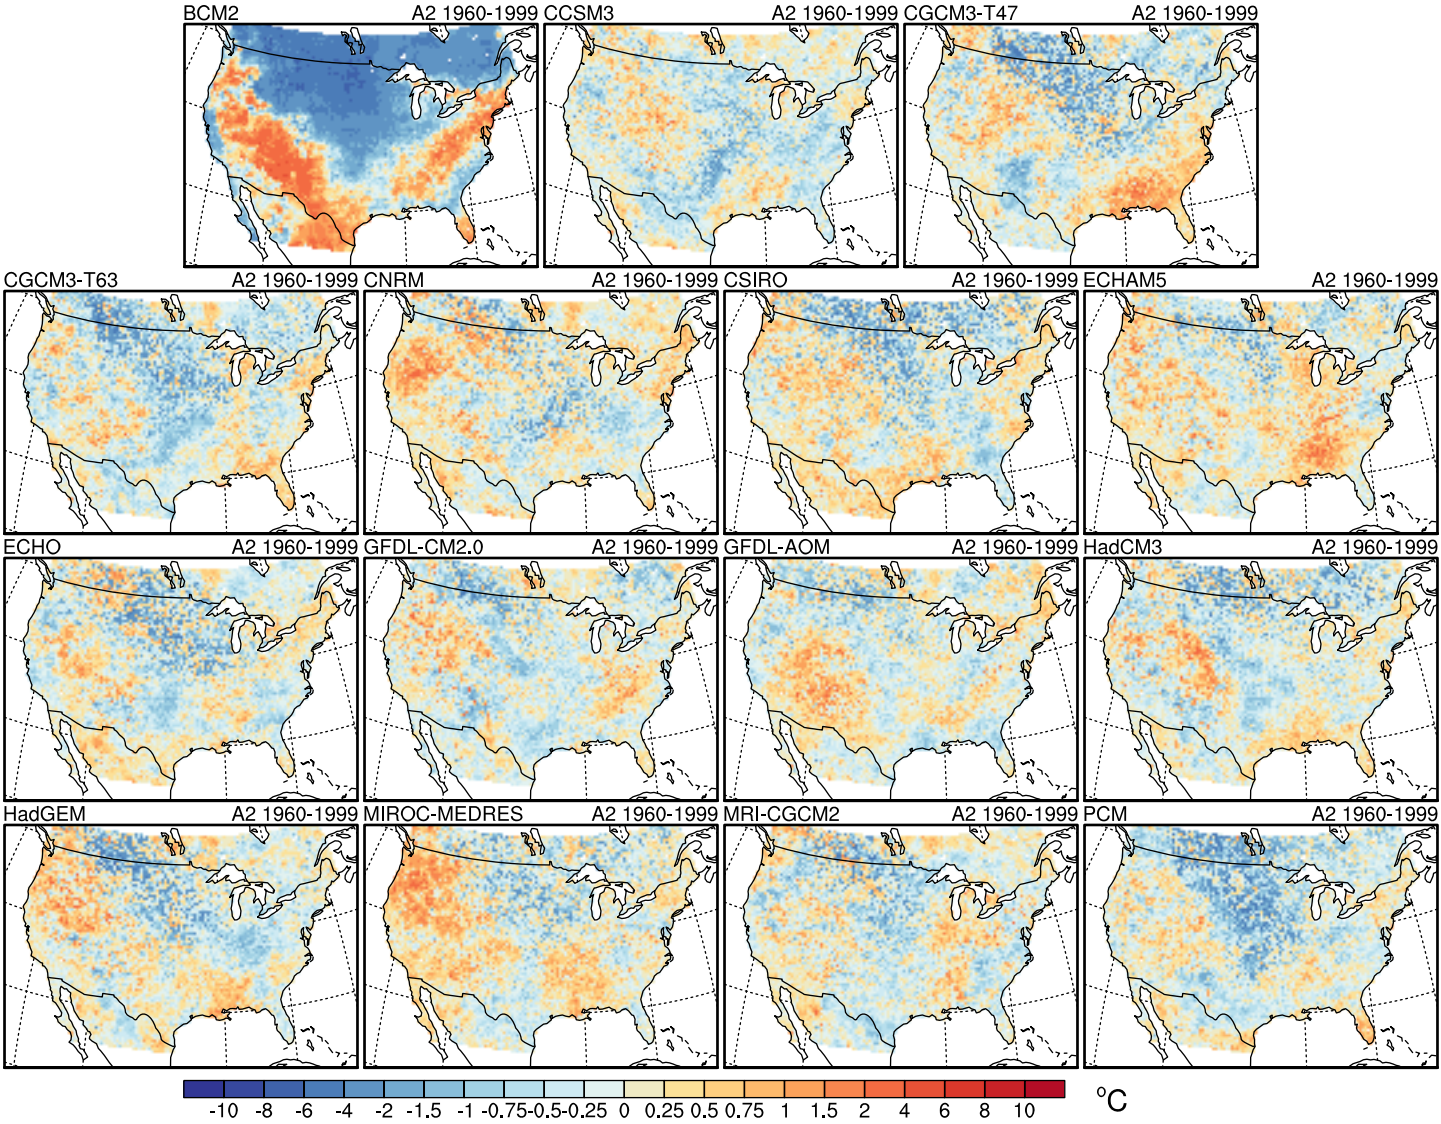
Tmax Bias in 99th Quantile



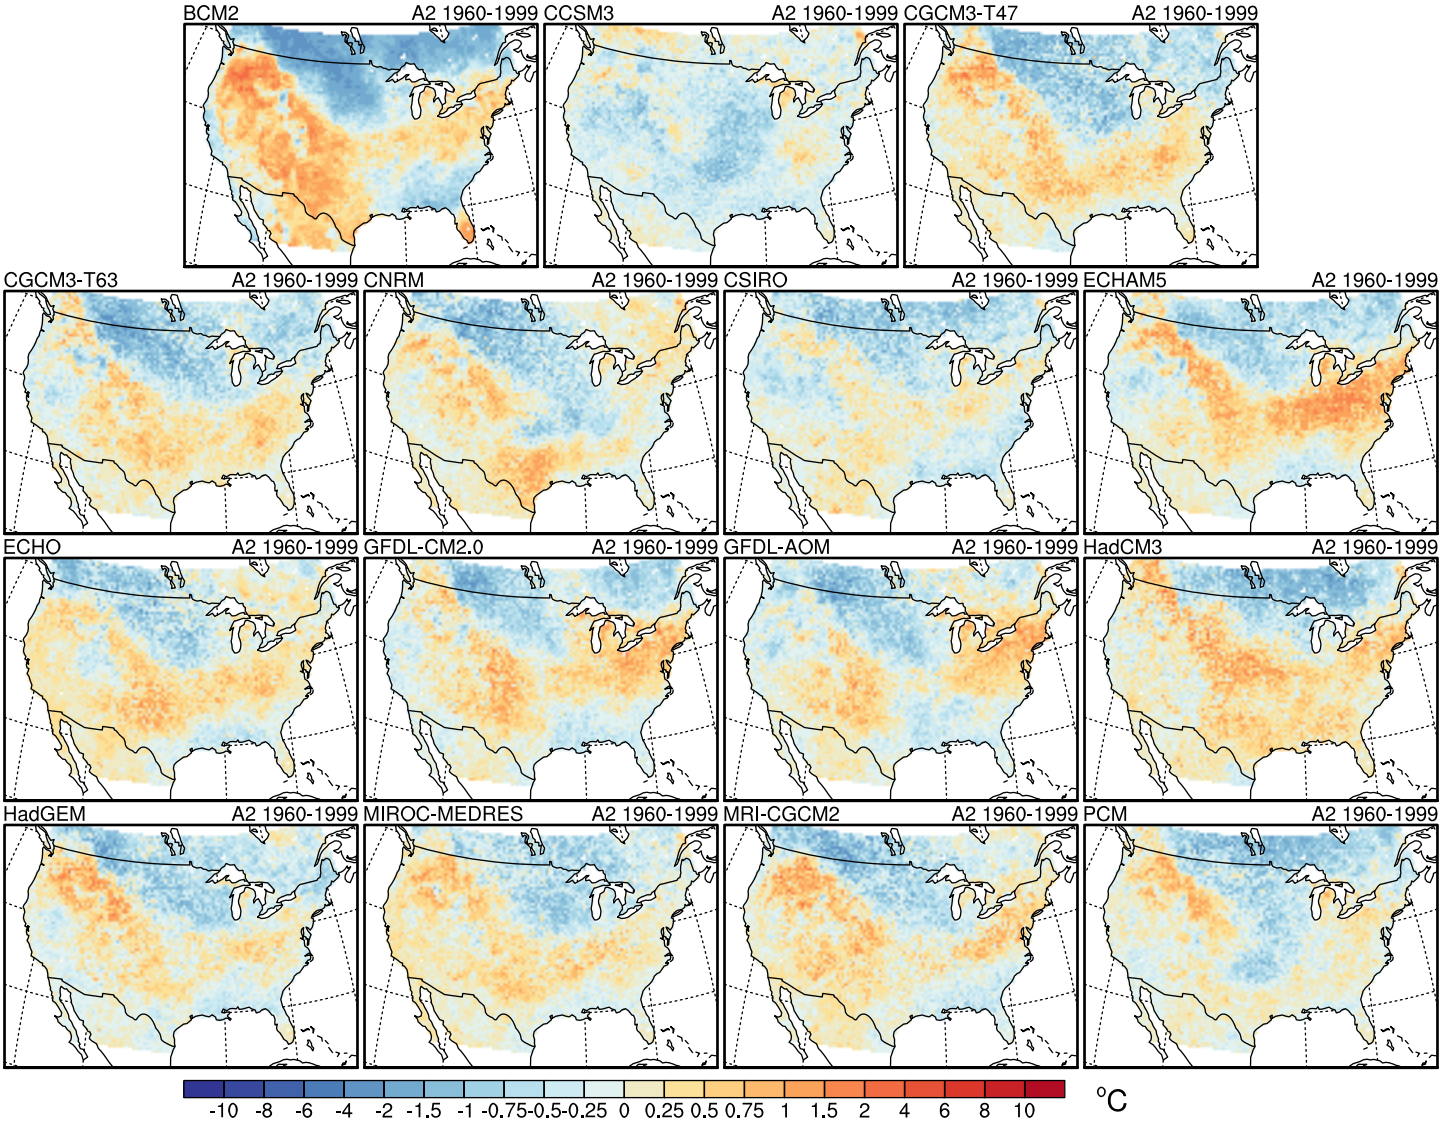
Tmax Bias in 99.9th Quantile



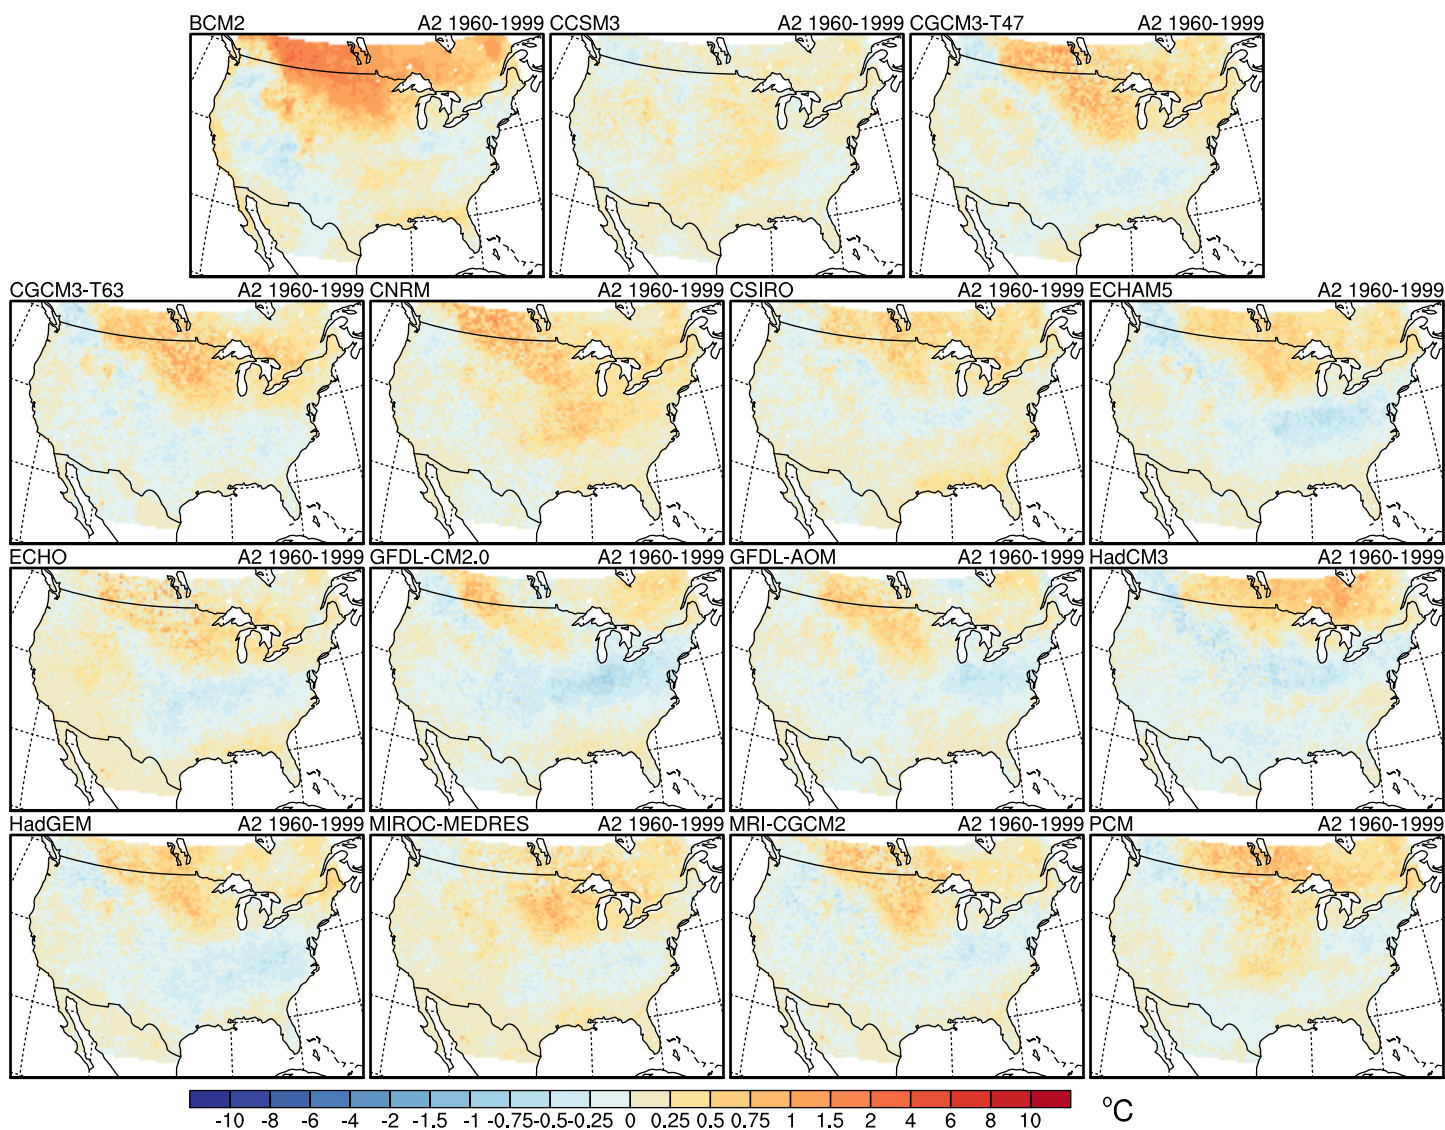
Tmin Bias in 0.1th Quantile



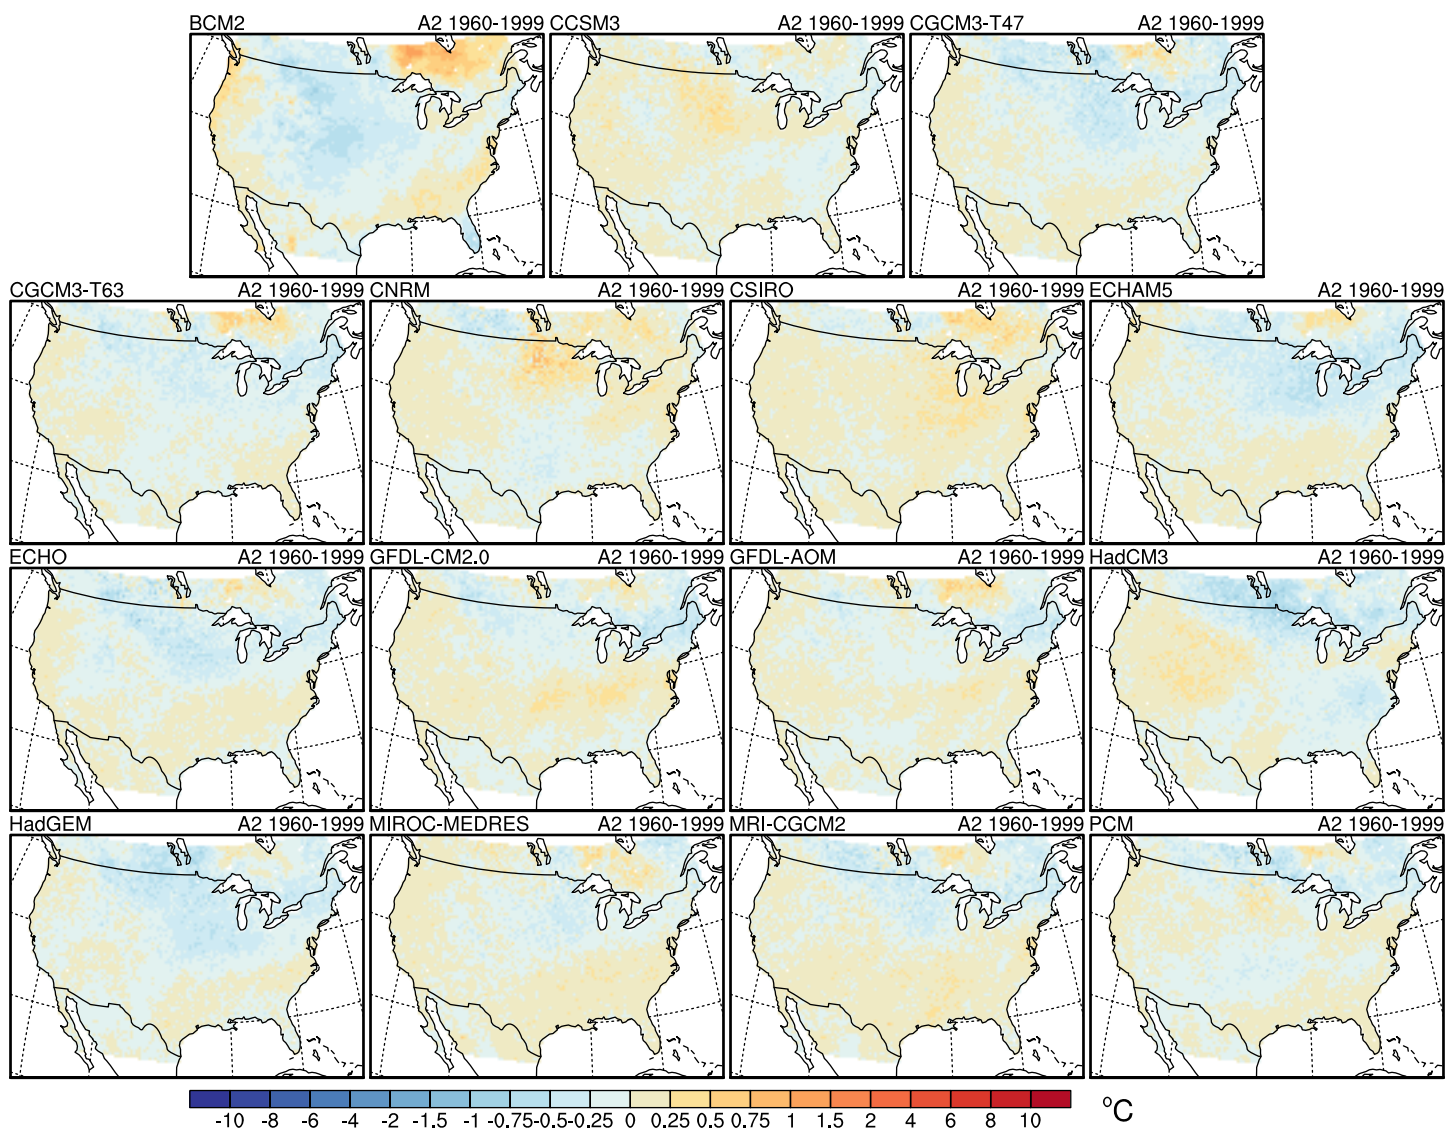
Tmin Bias in 1st Quantile



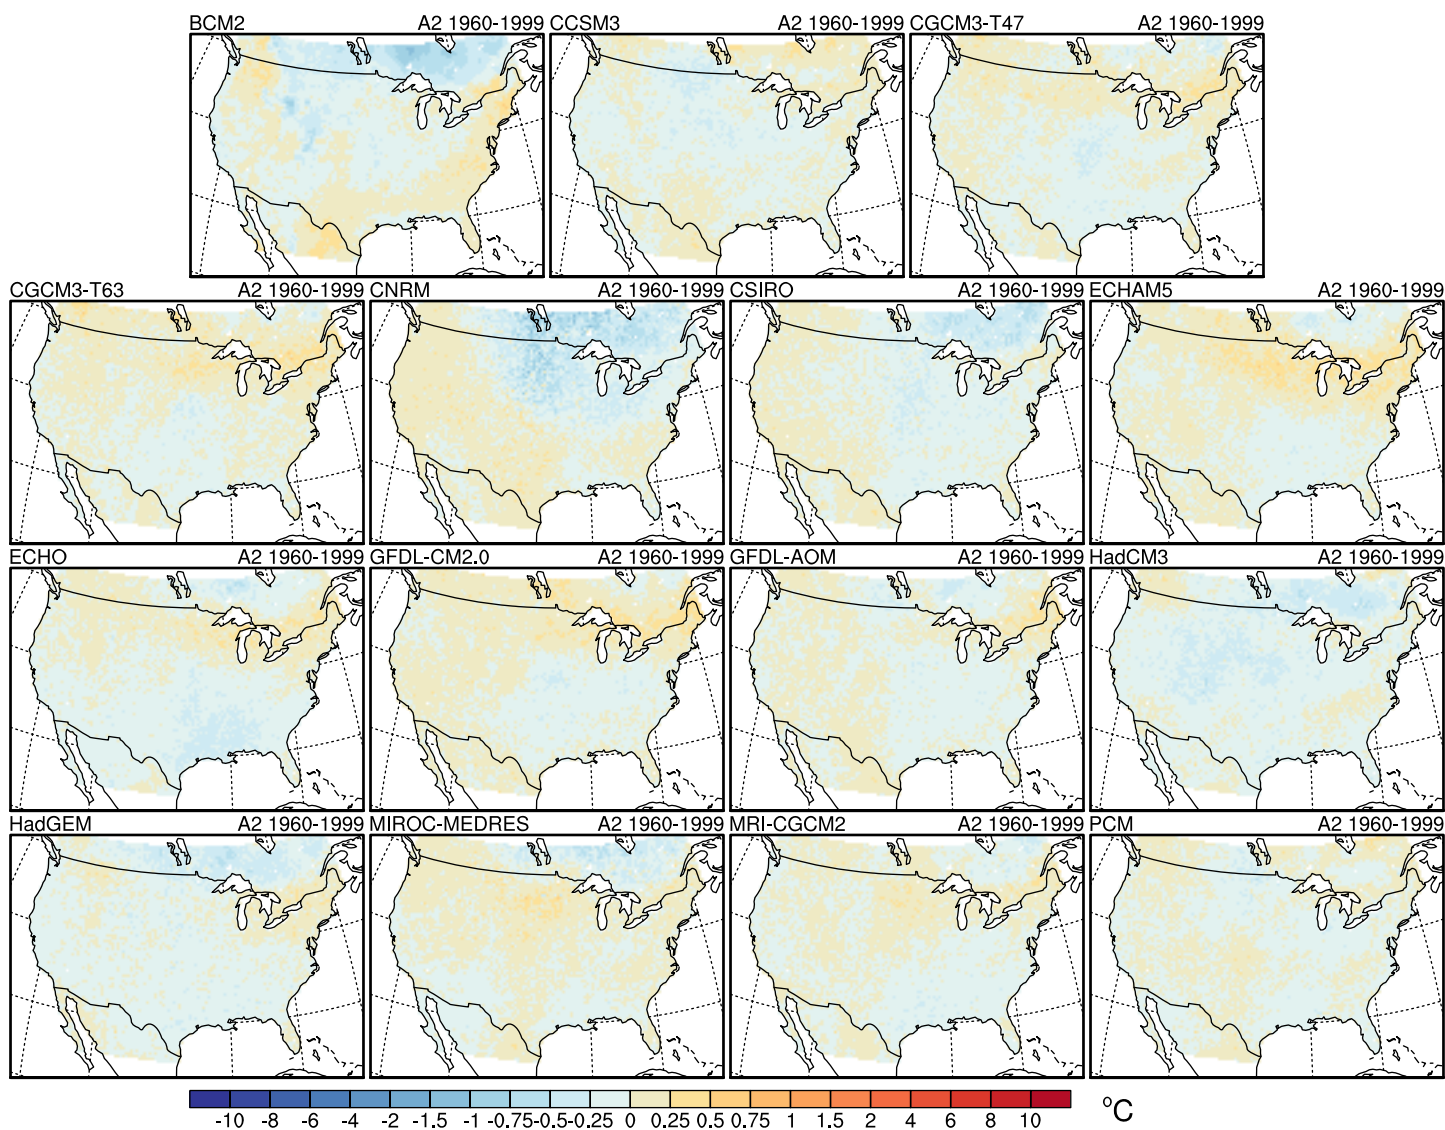
Tmin Bias in 10th Quantile



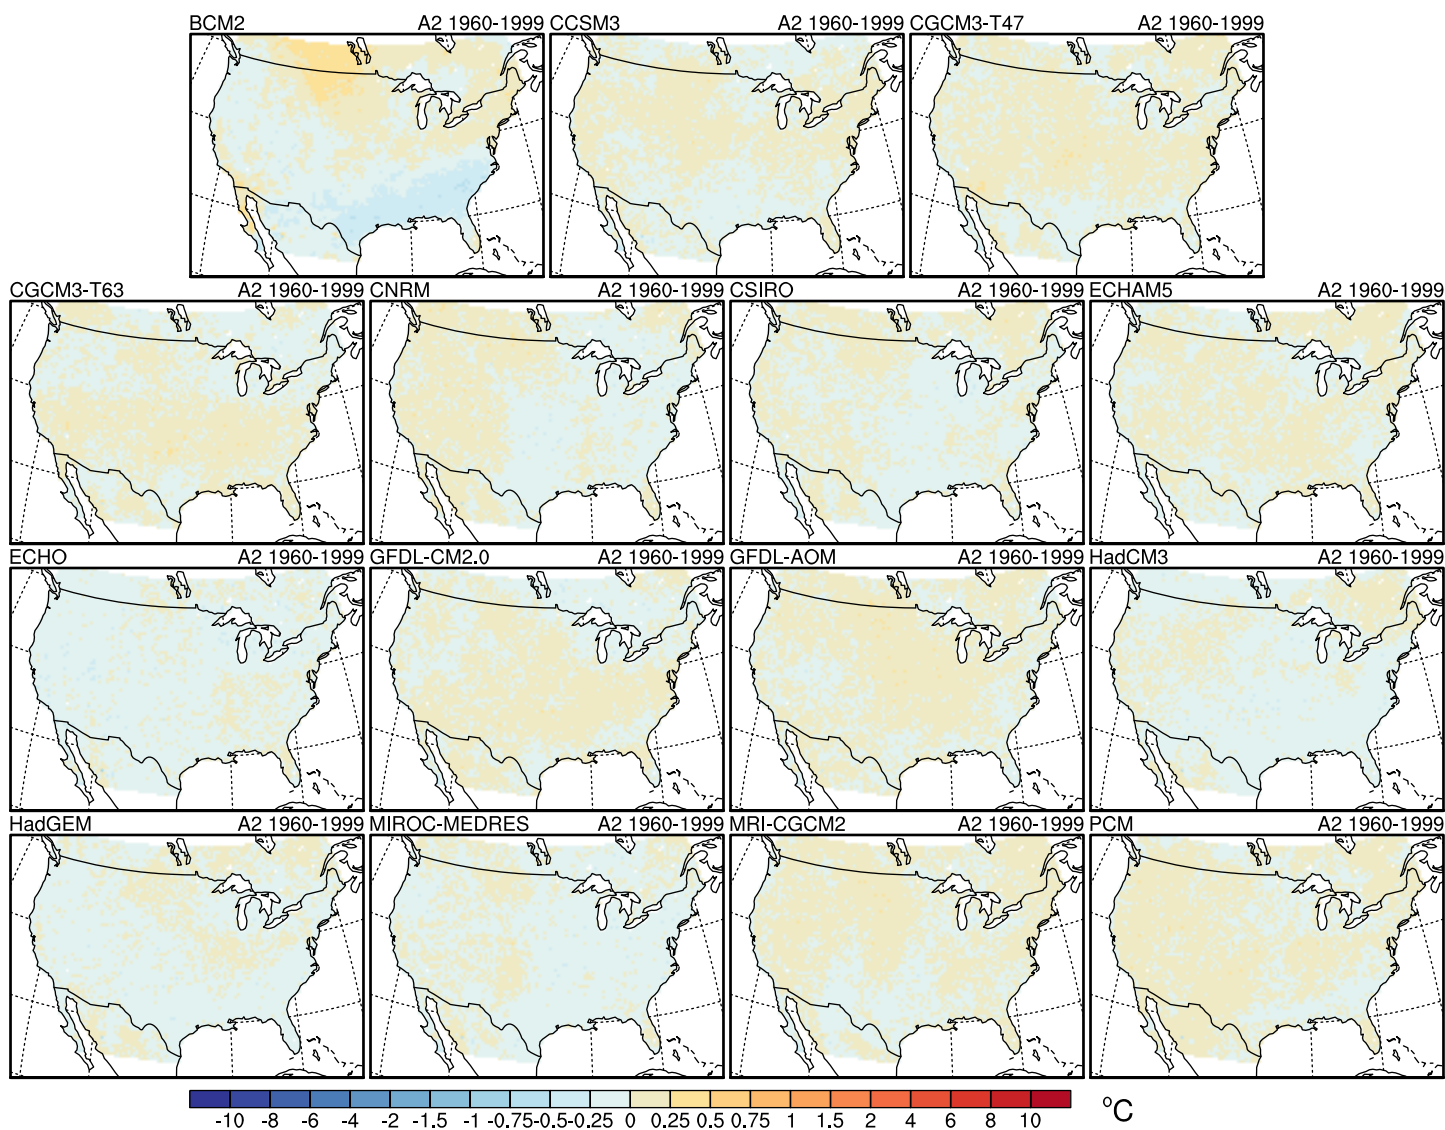
Tmin Bias in 25th Quantile



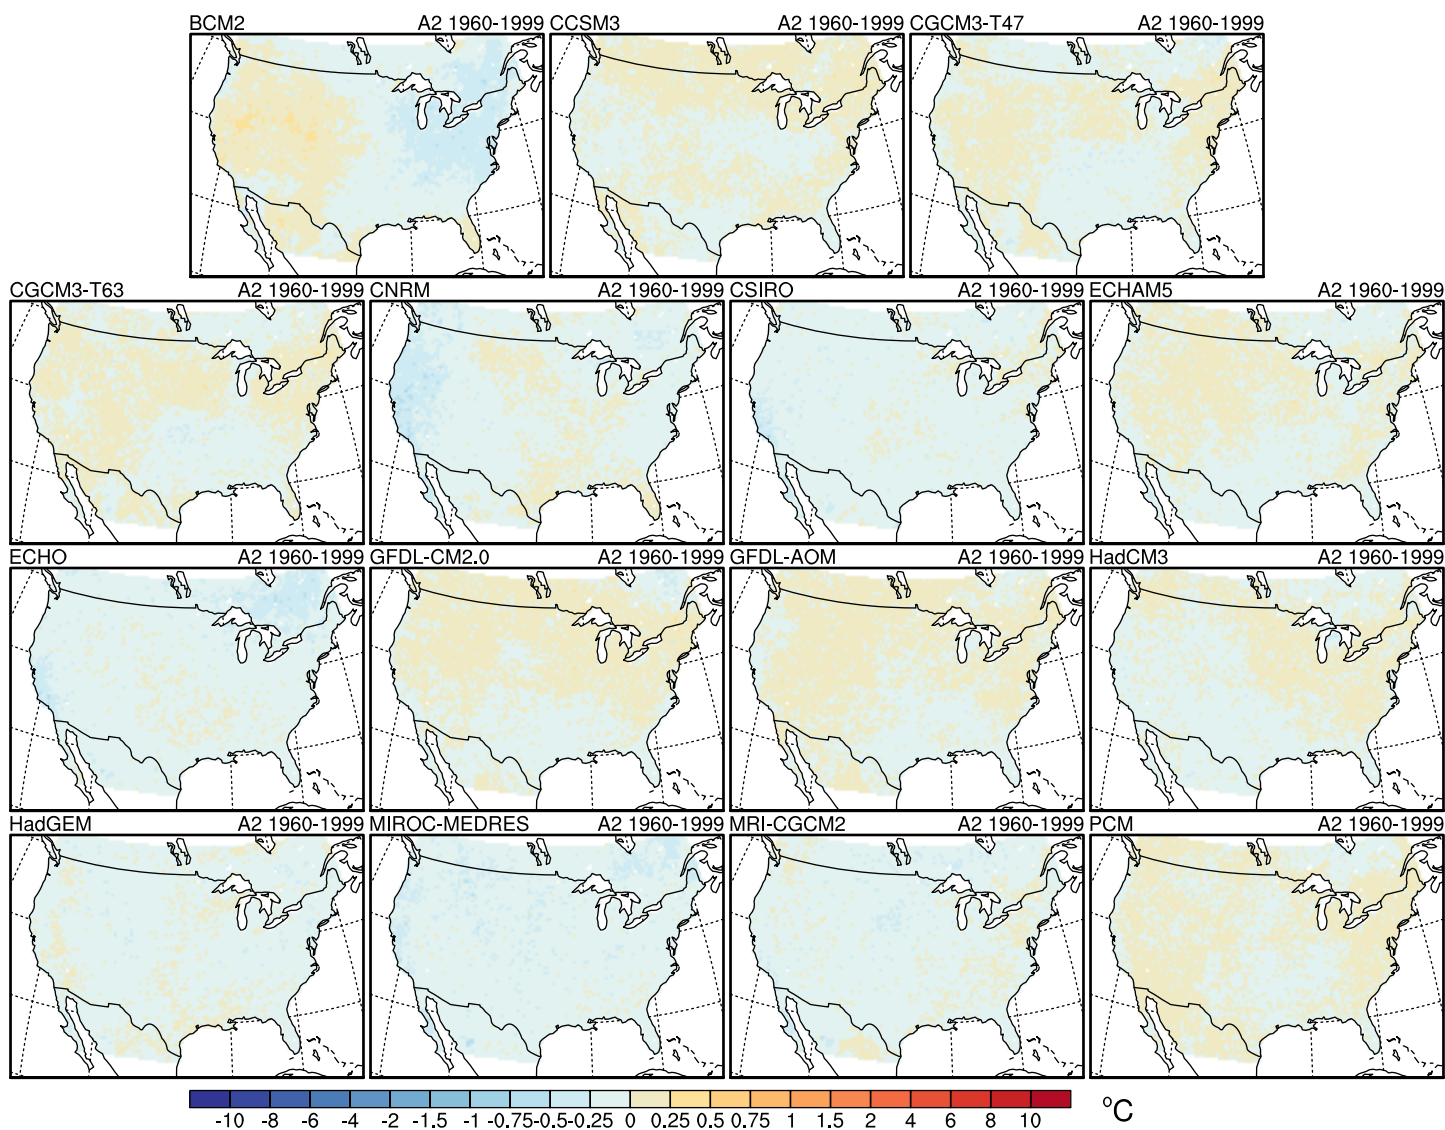
Tmin Bias in 50th Quantile



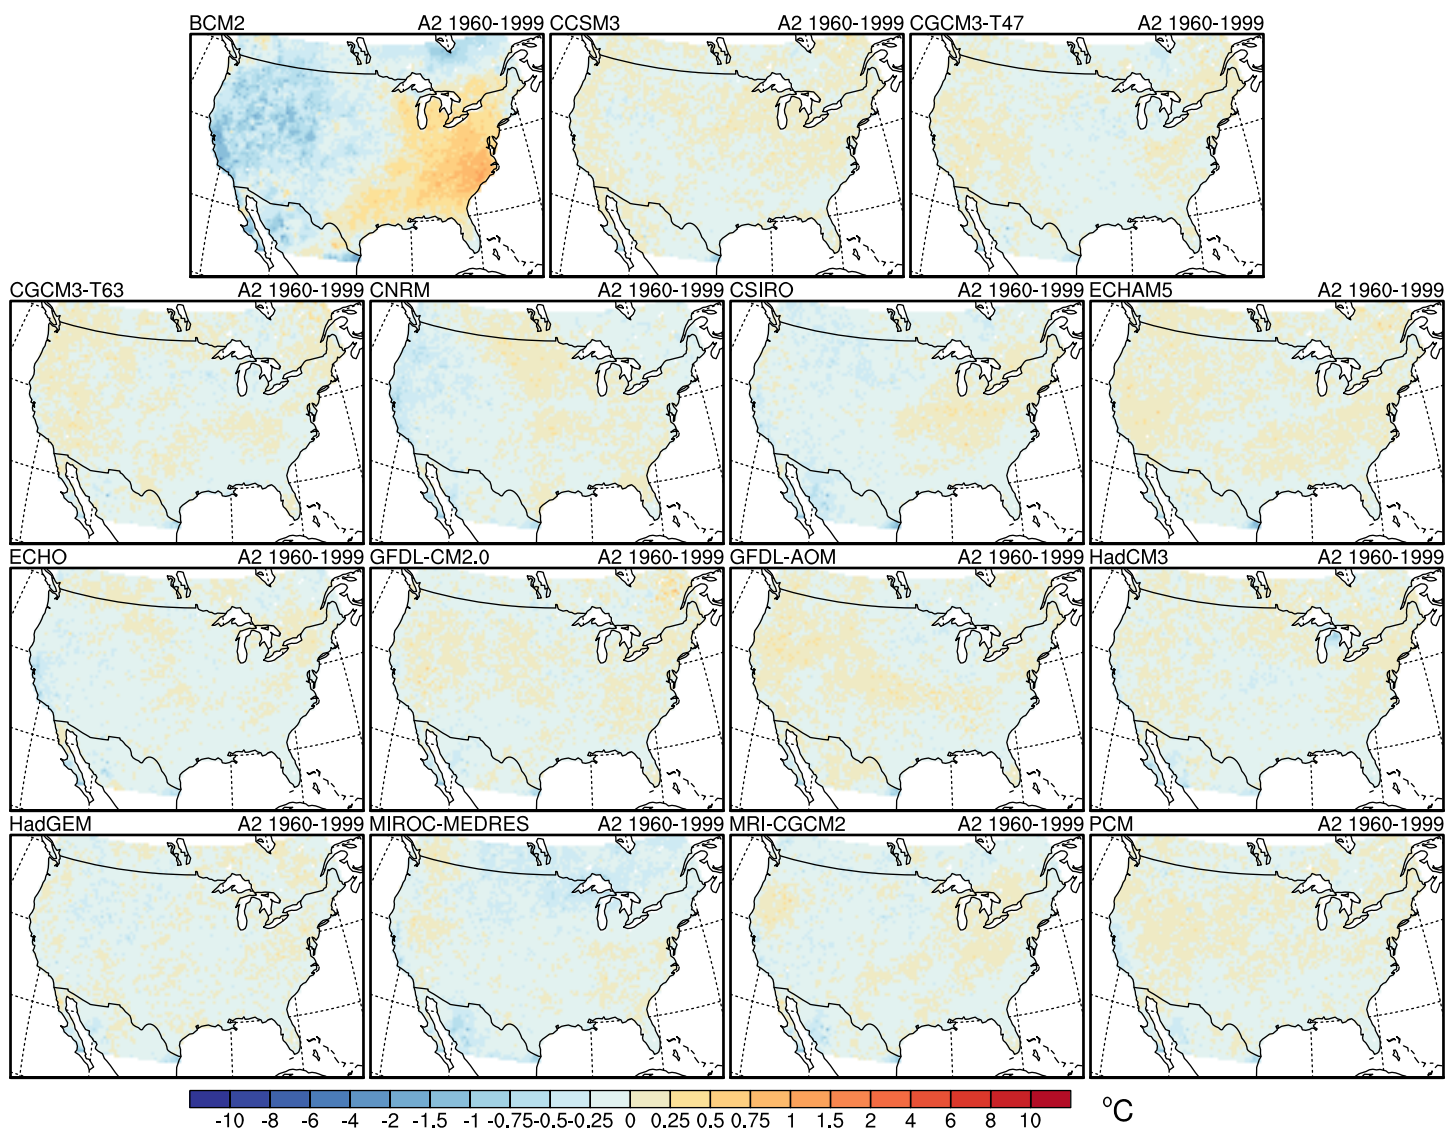
Tmin Bias in 75th Quantile



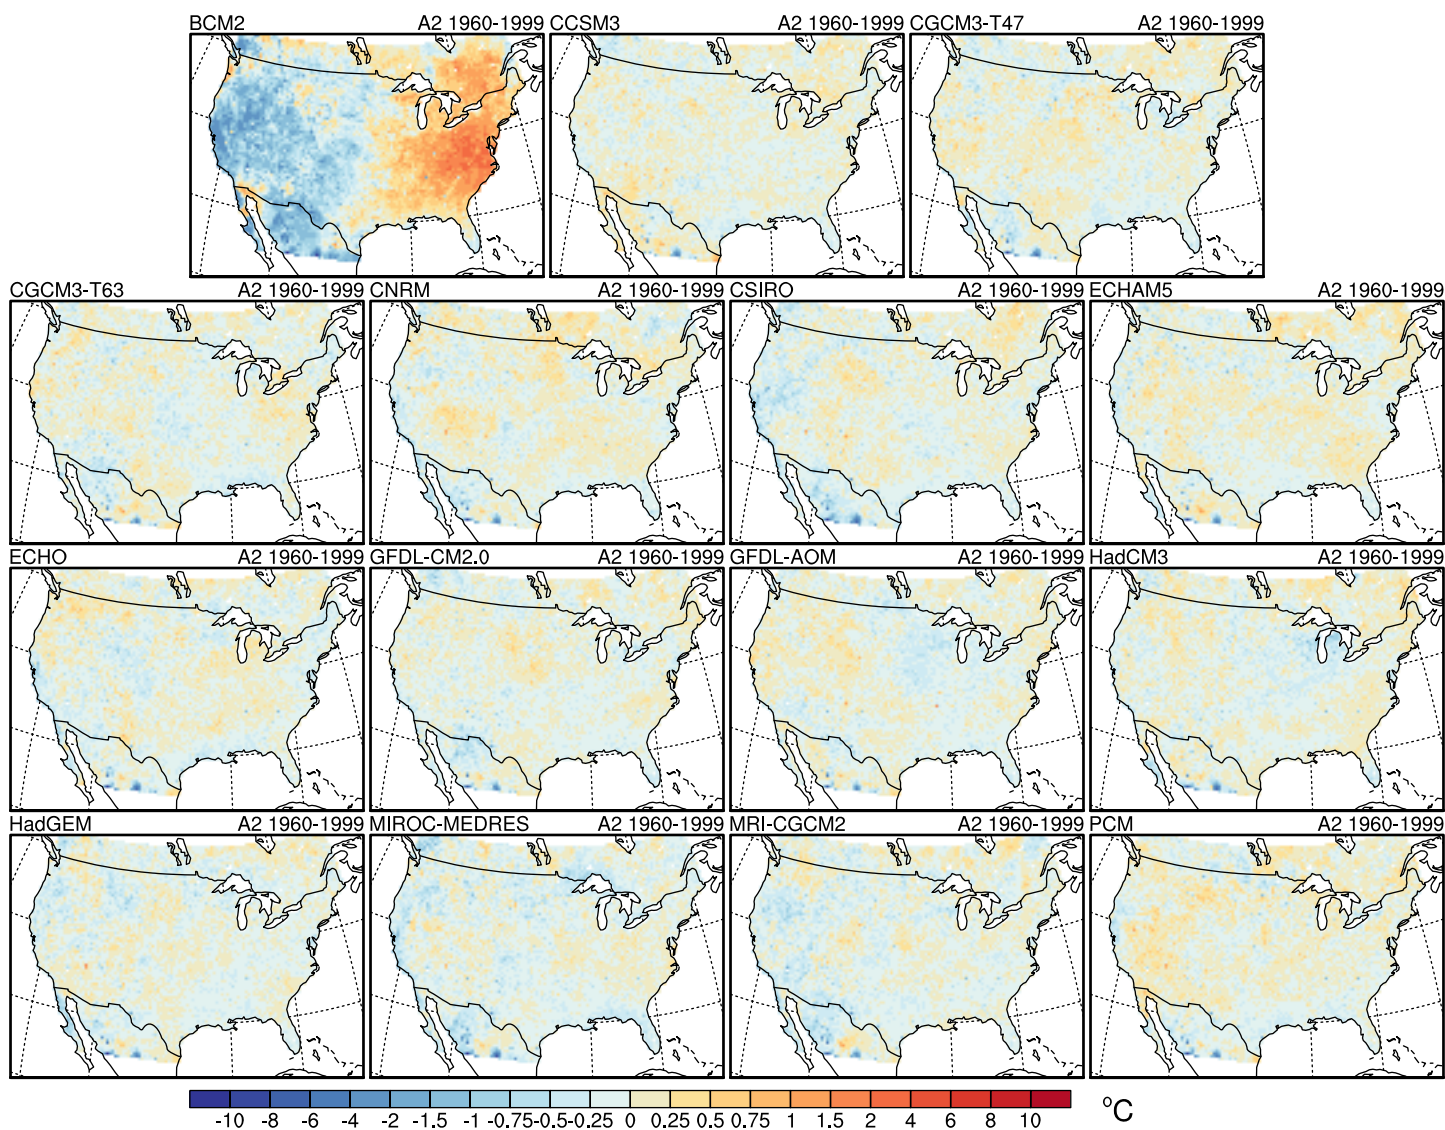
Tmin Bias in 90th Quantile



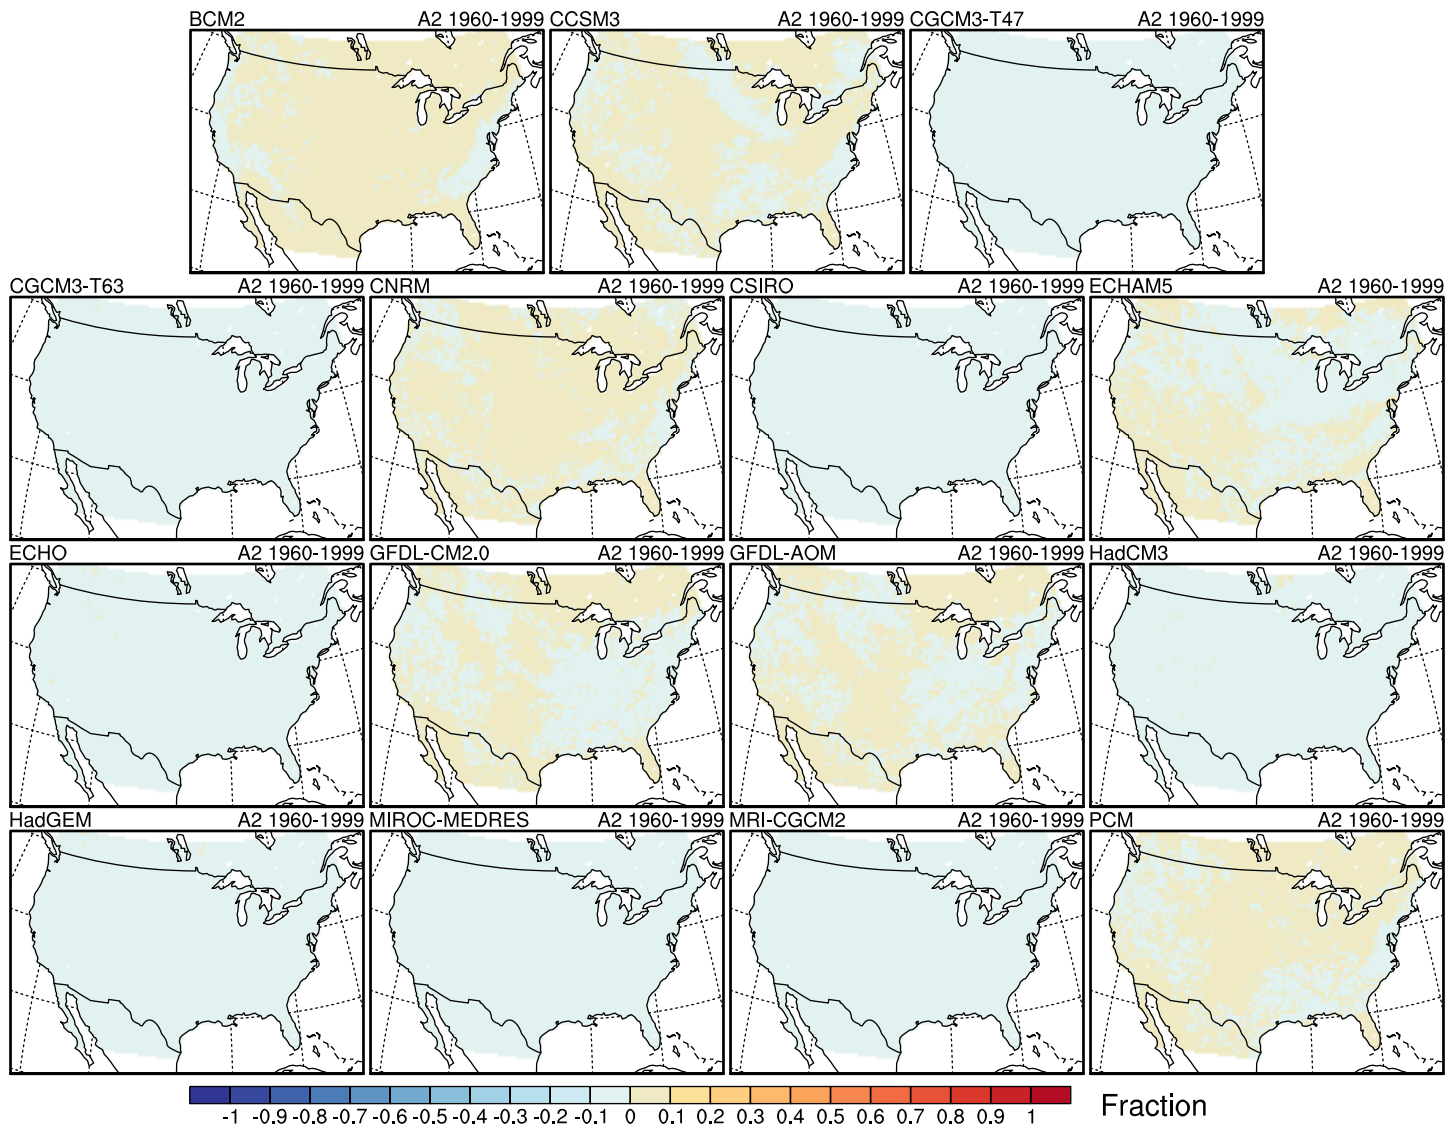
Tmin Bias in 99th Quantile



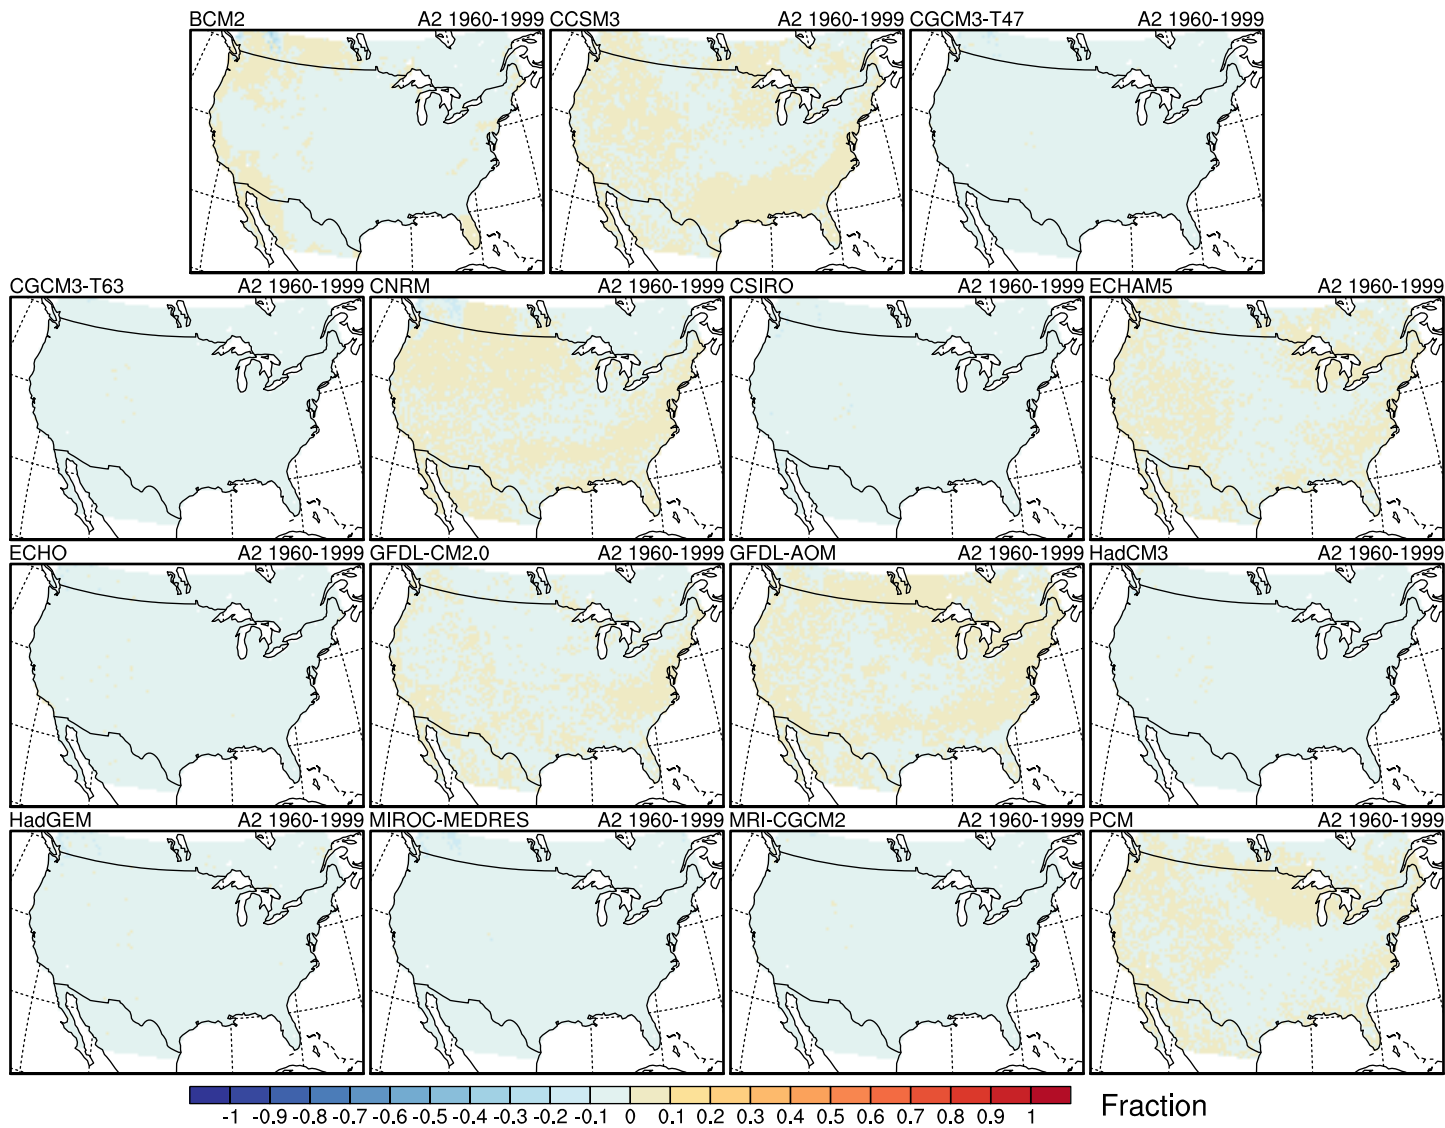
Tmin Bias in 99.9th Quantile



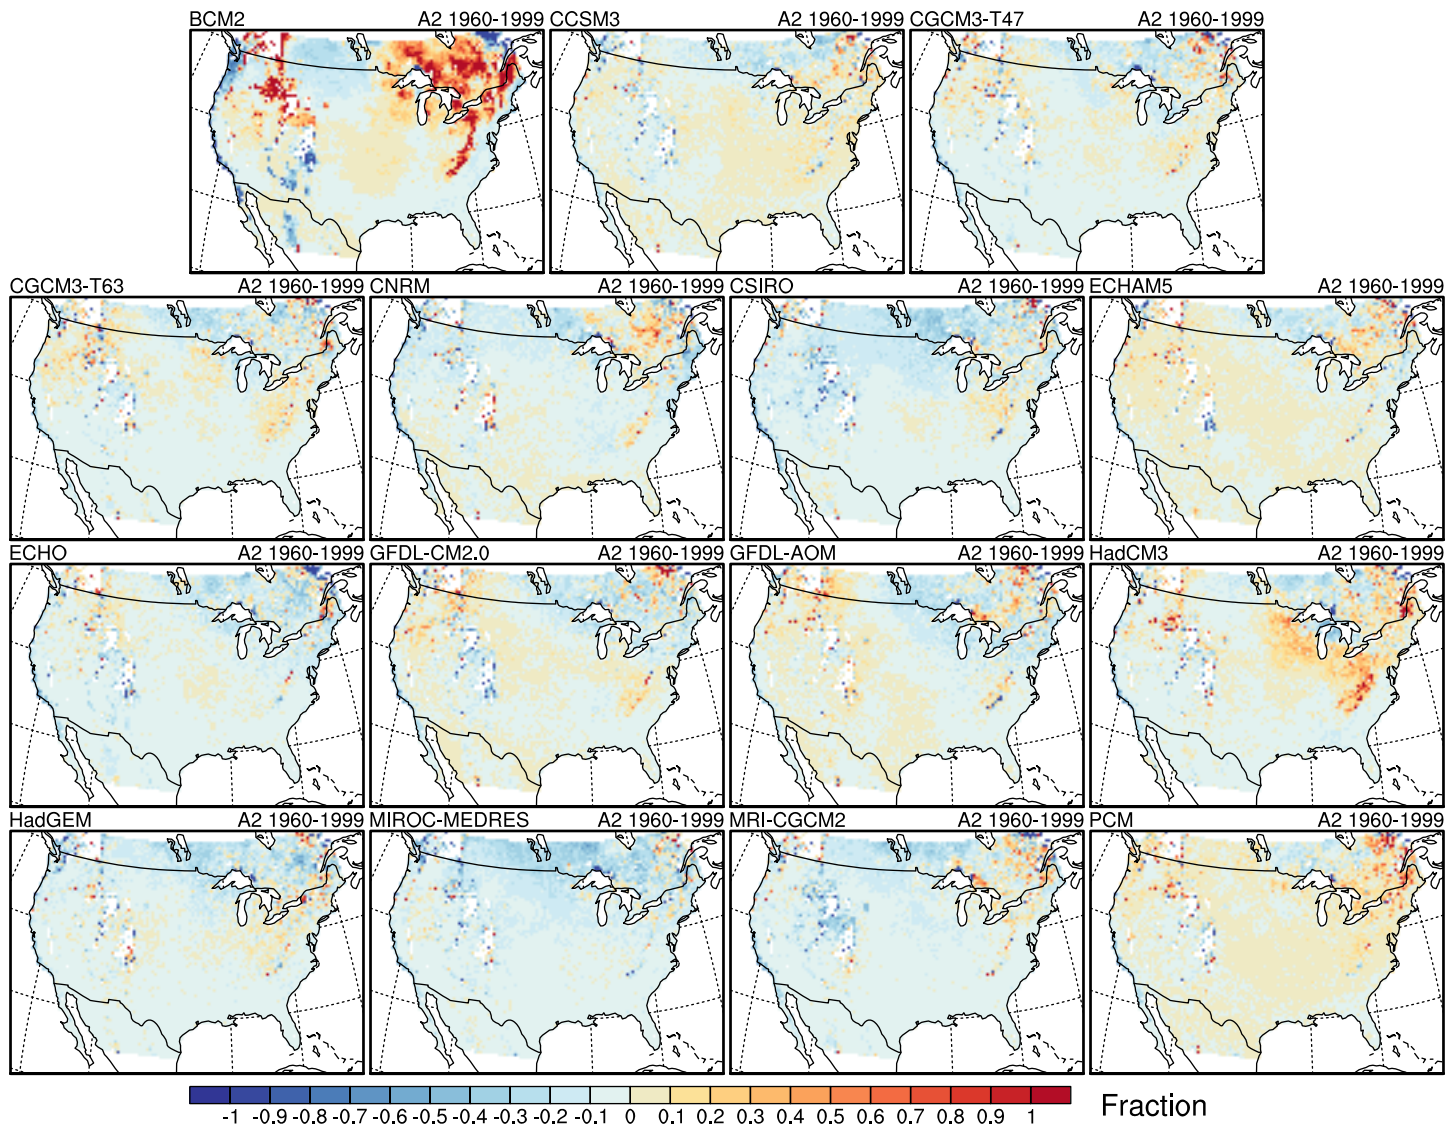
Bias in number of days with Tmax > 32 F



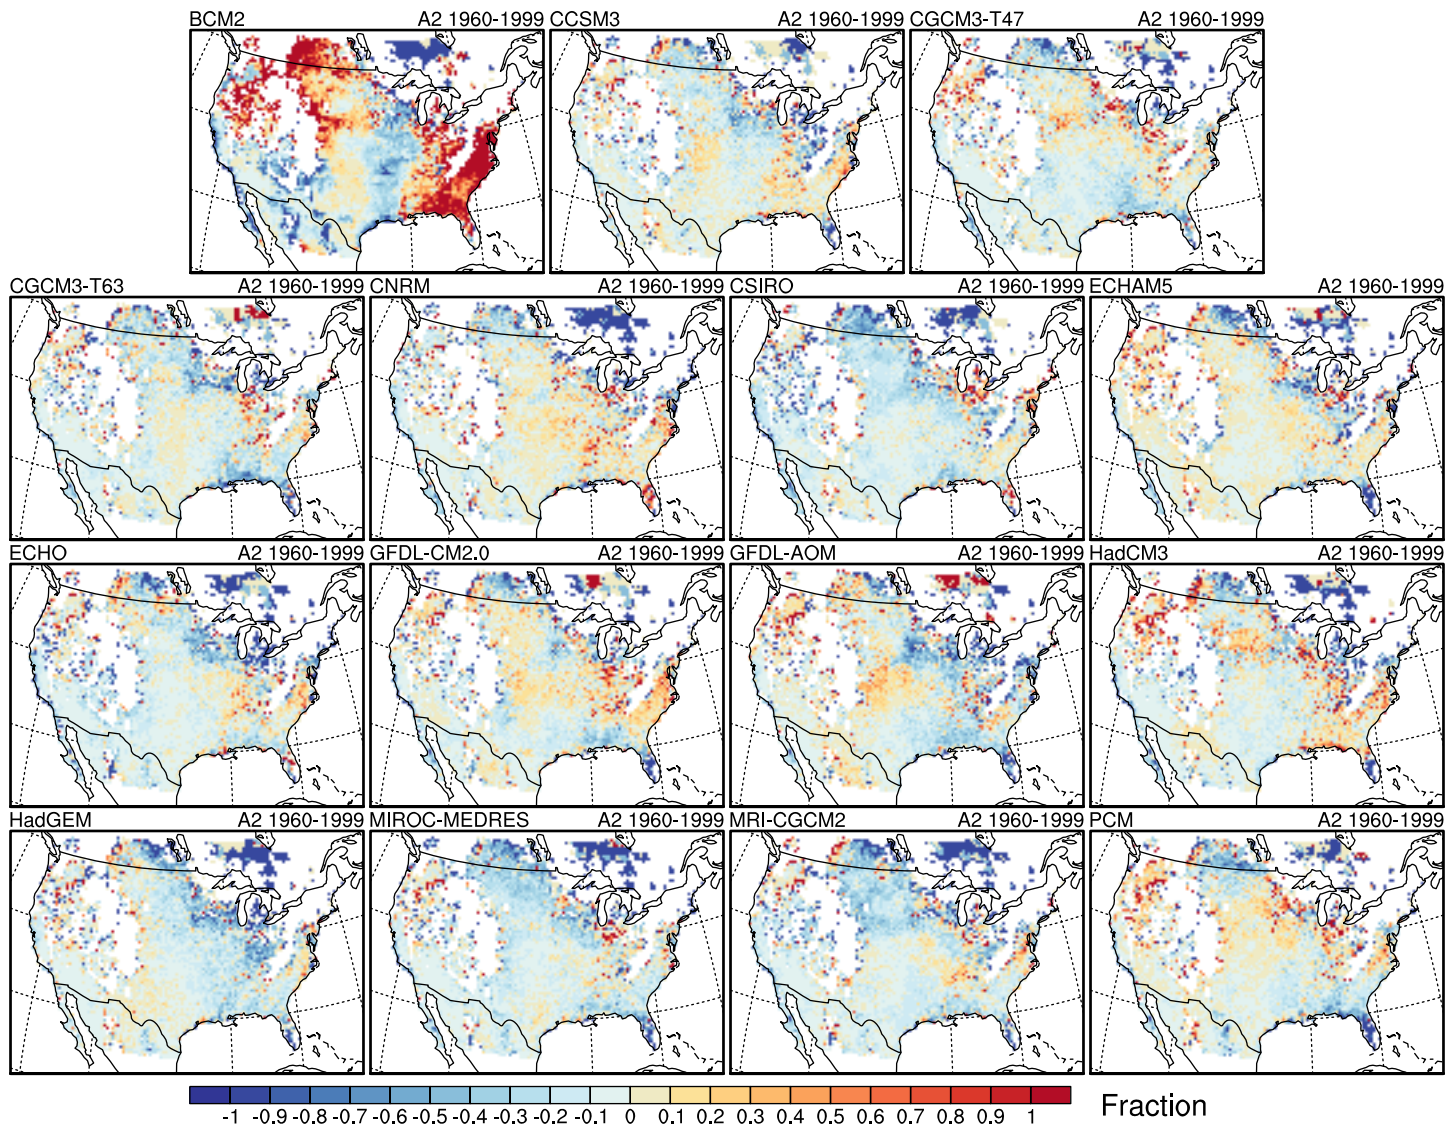
Bias in number of days with Tmax > 65 F



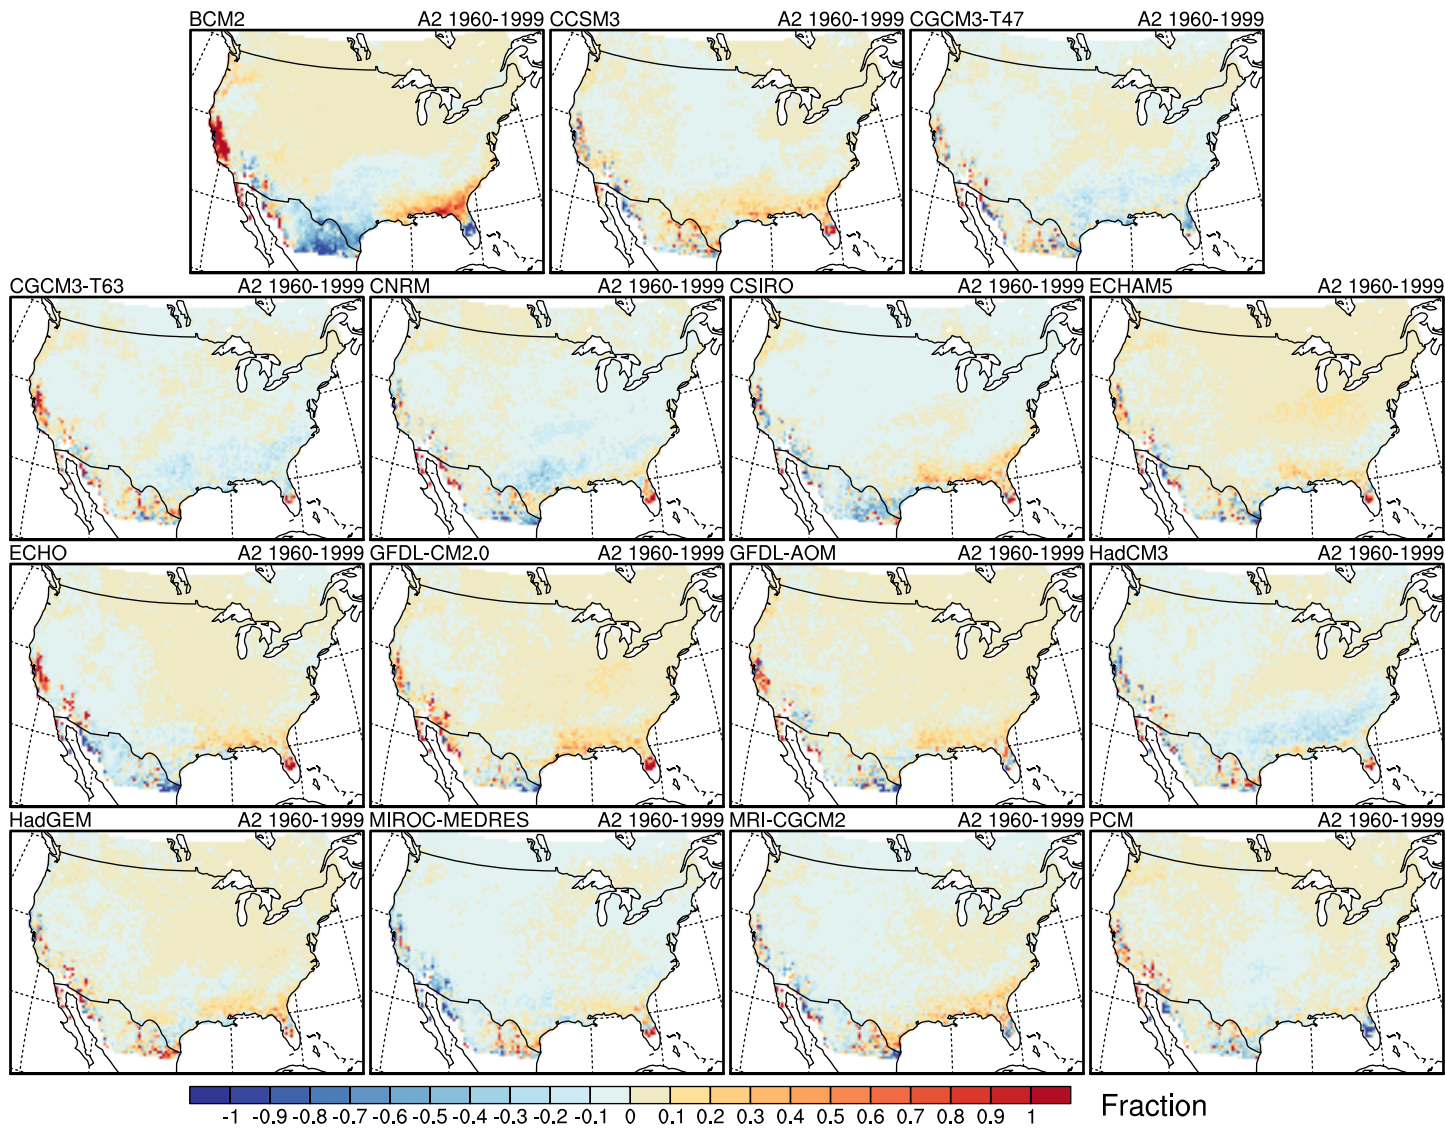
Bias in number of days with Tmax > 90 F



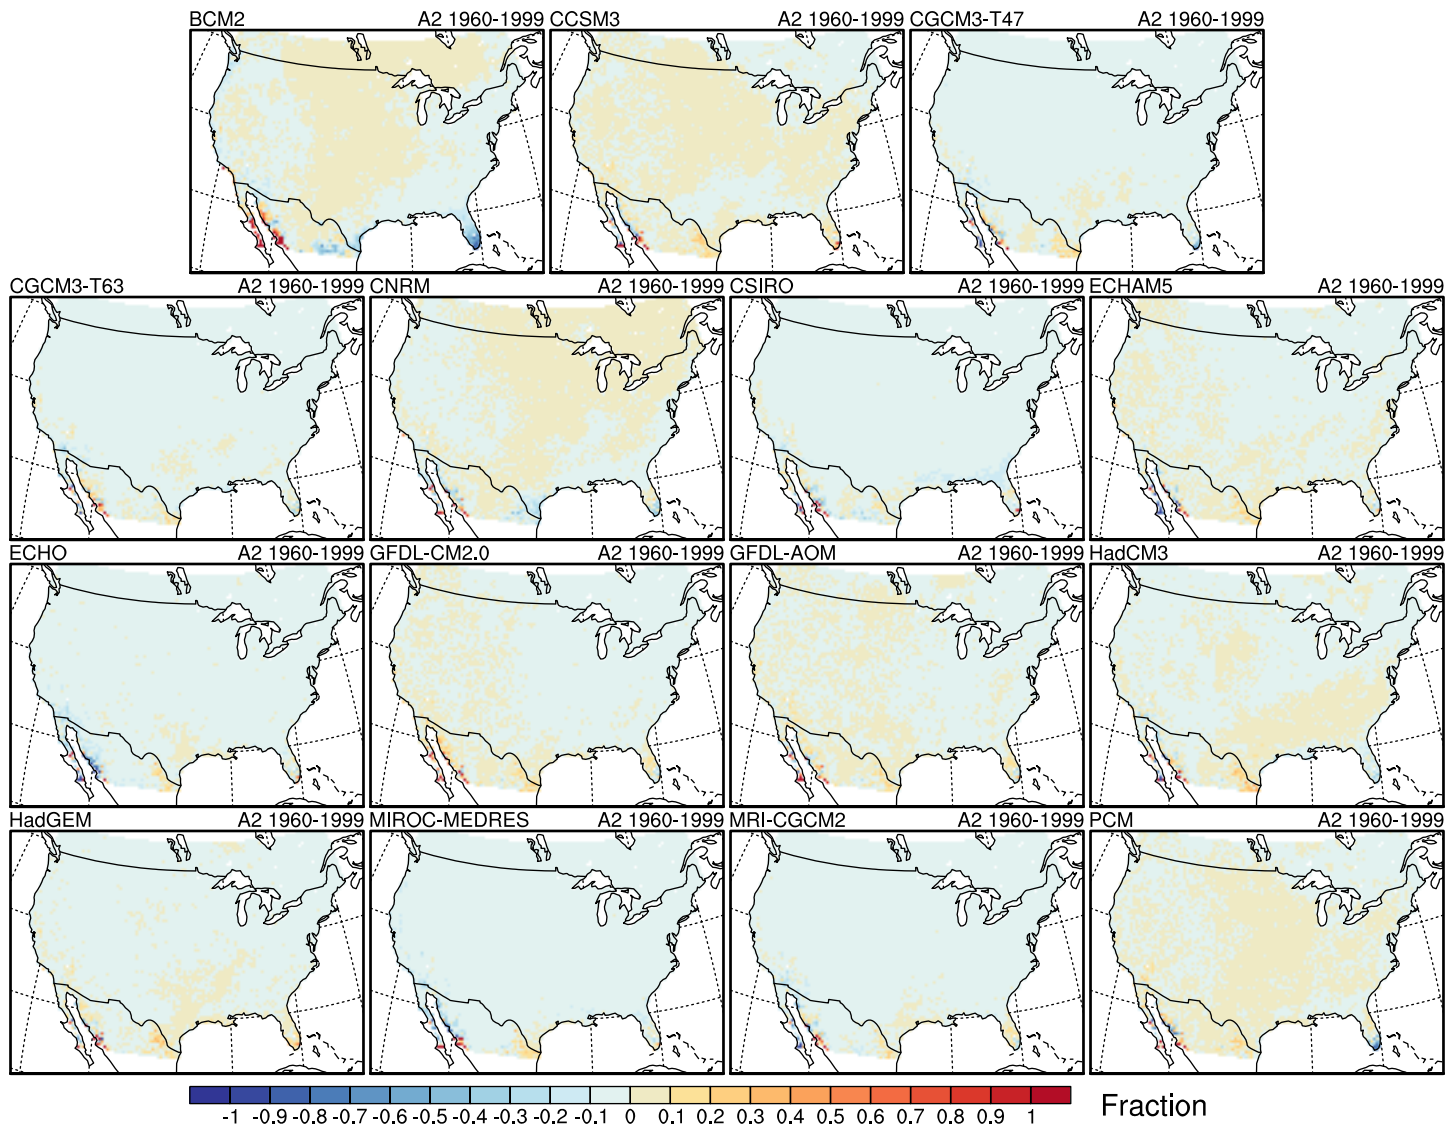
Bias in number of days with Tmax > 100 F



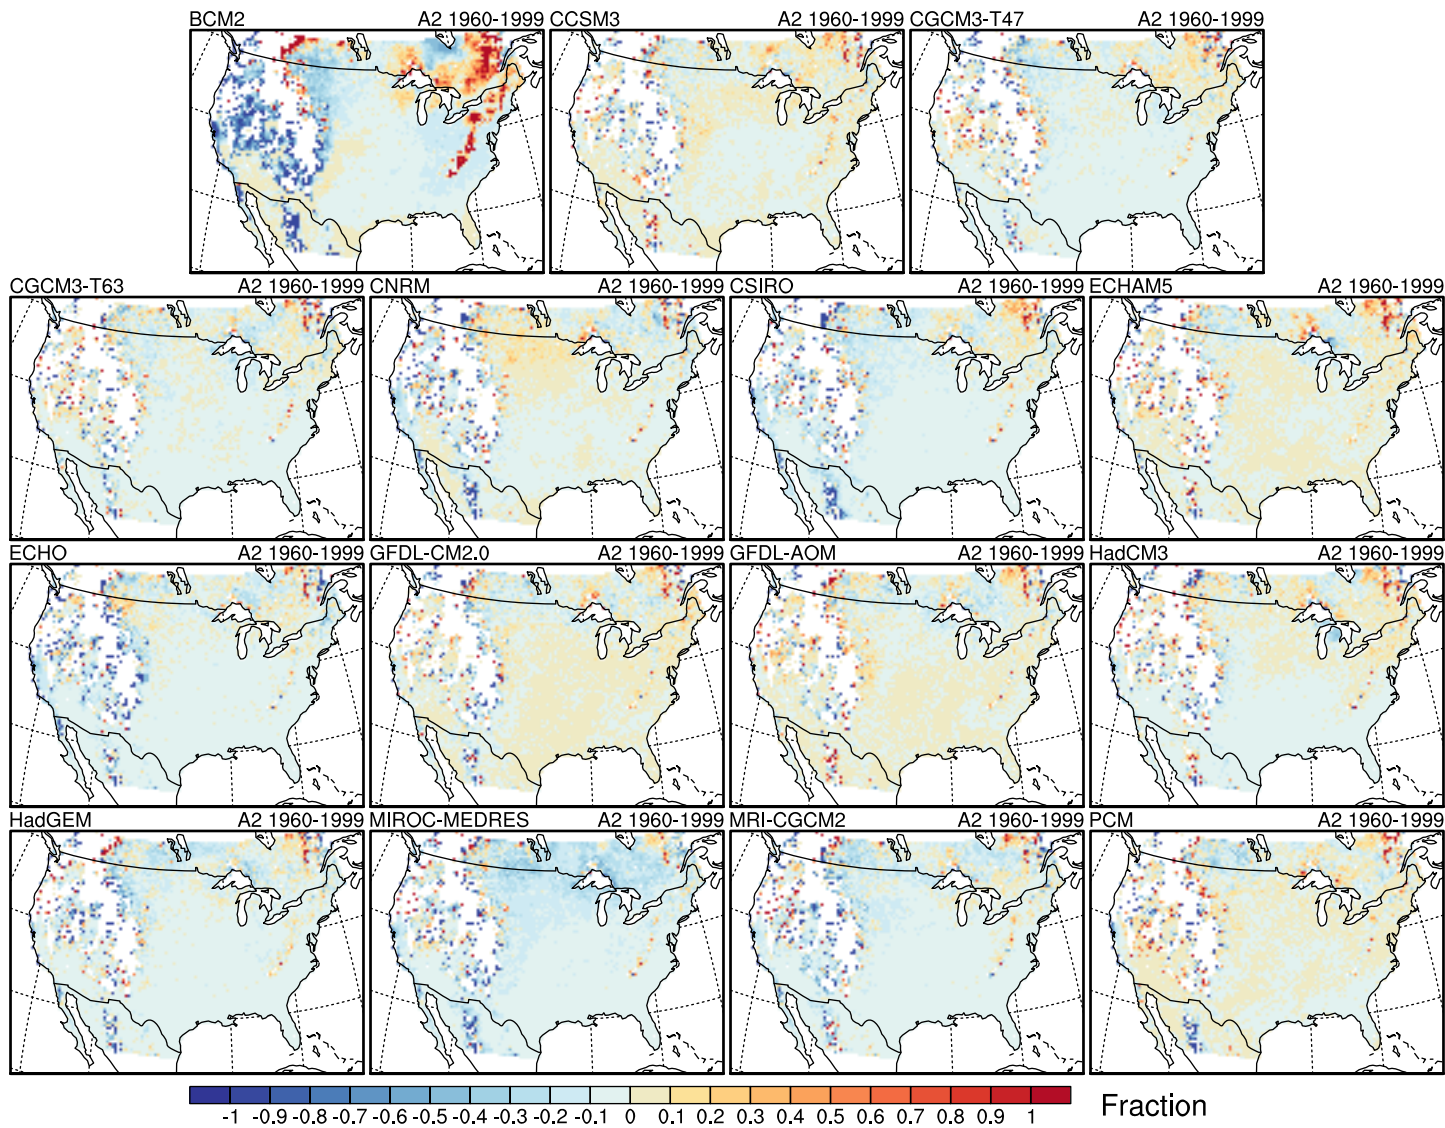
Bias in number of days with $T_{min} < 20\text{ F}$



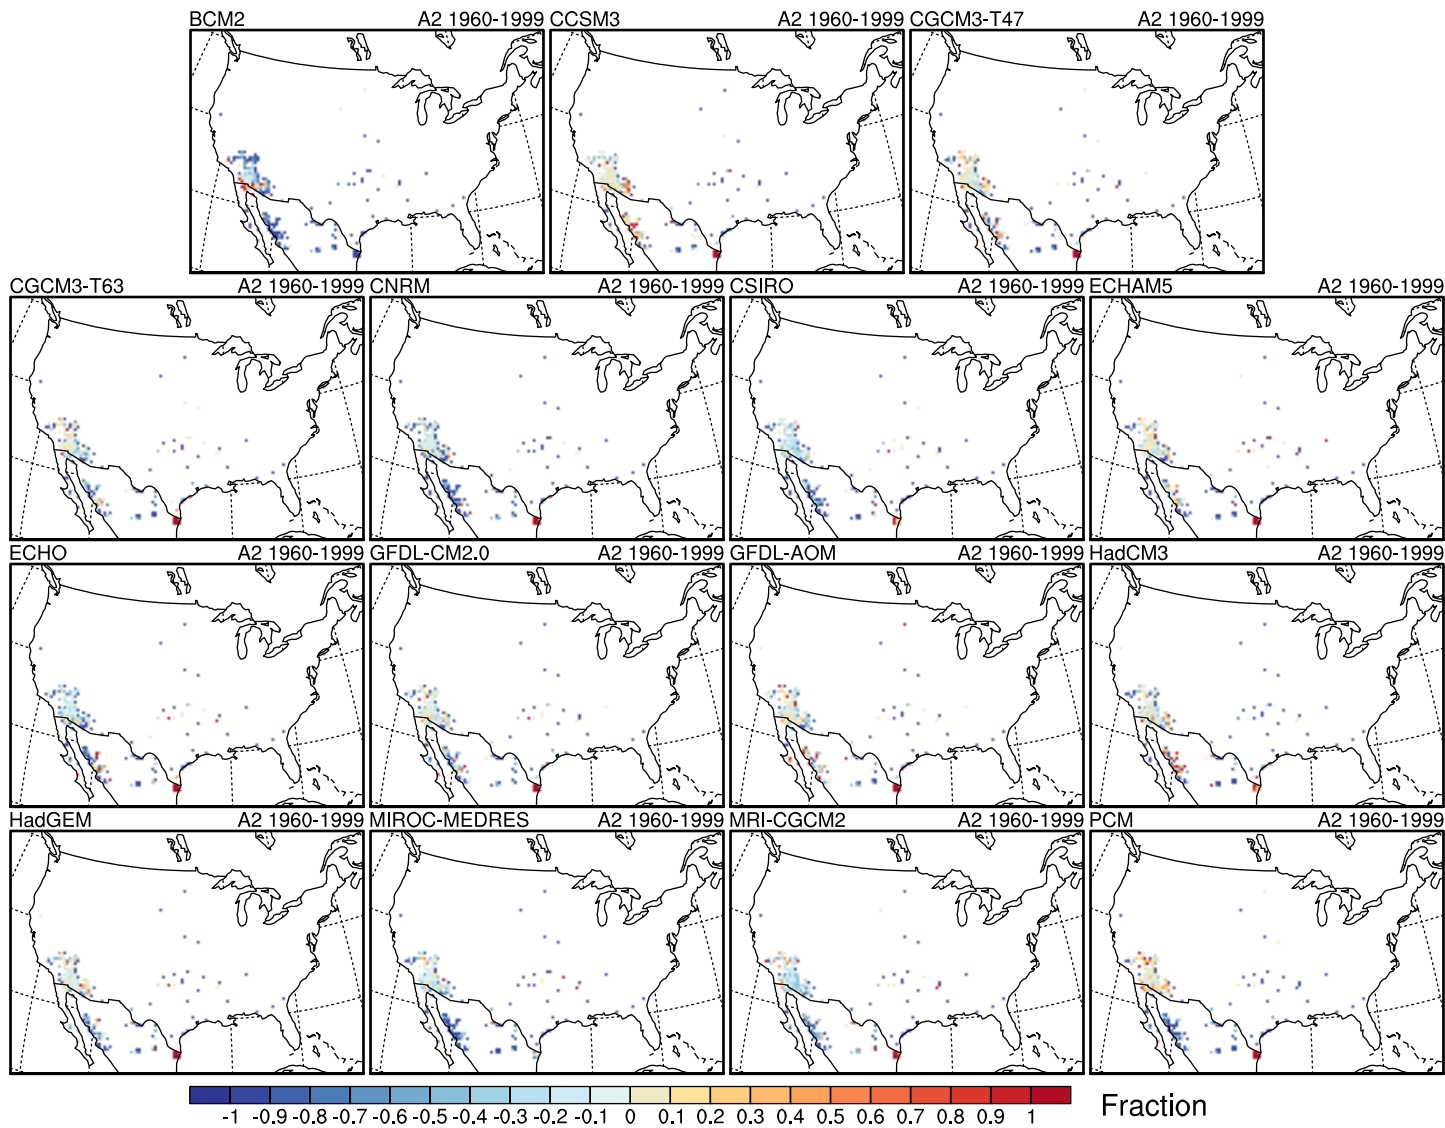
Bias in number of days with $T_{min} < 32\text{ F}$



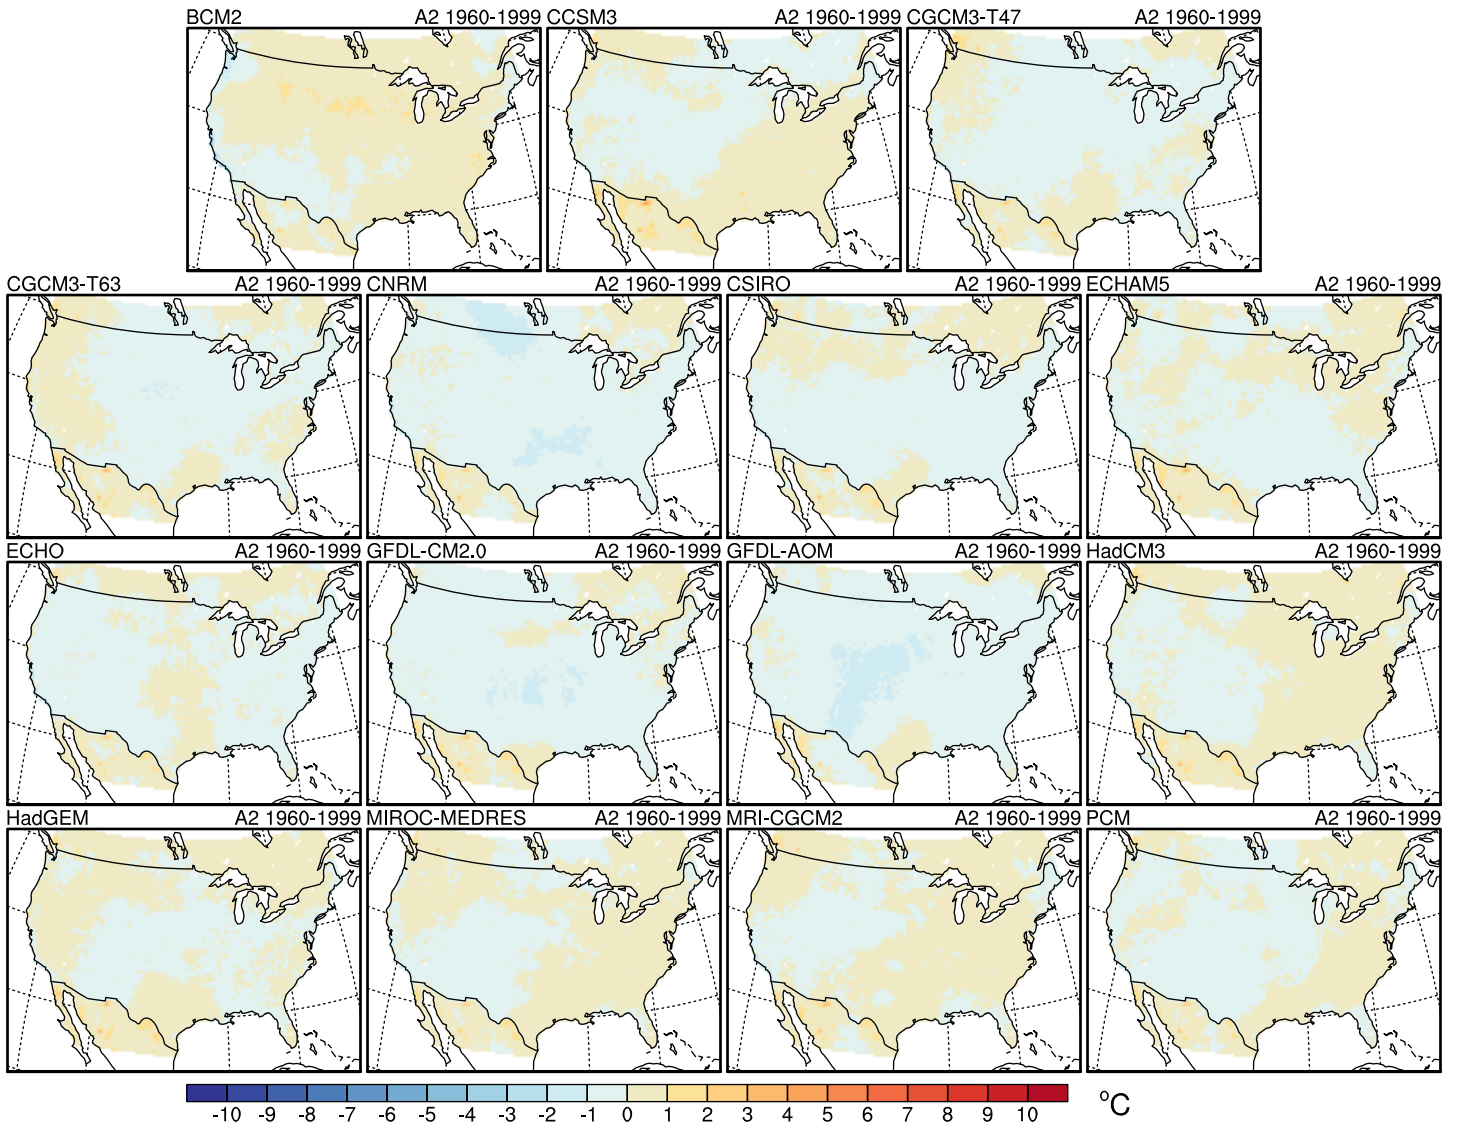
Bias in number of days with $T_{min} < 65\text{ F}$



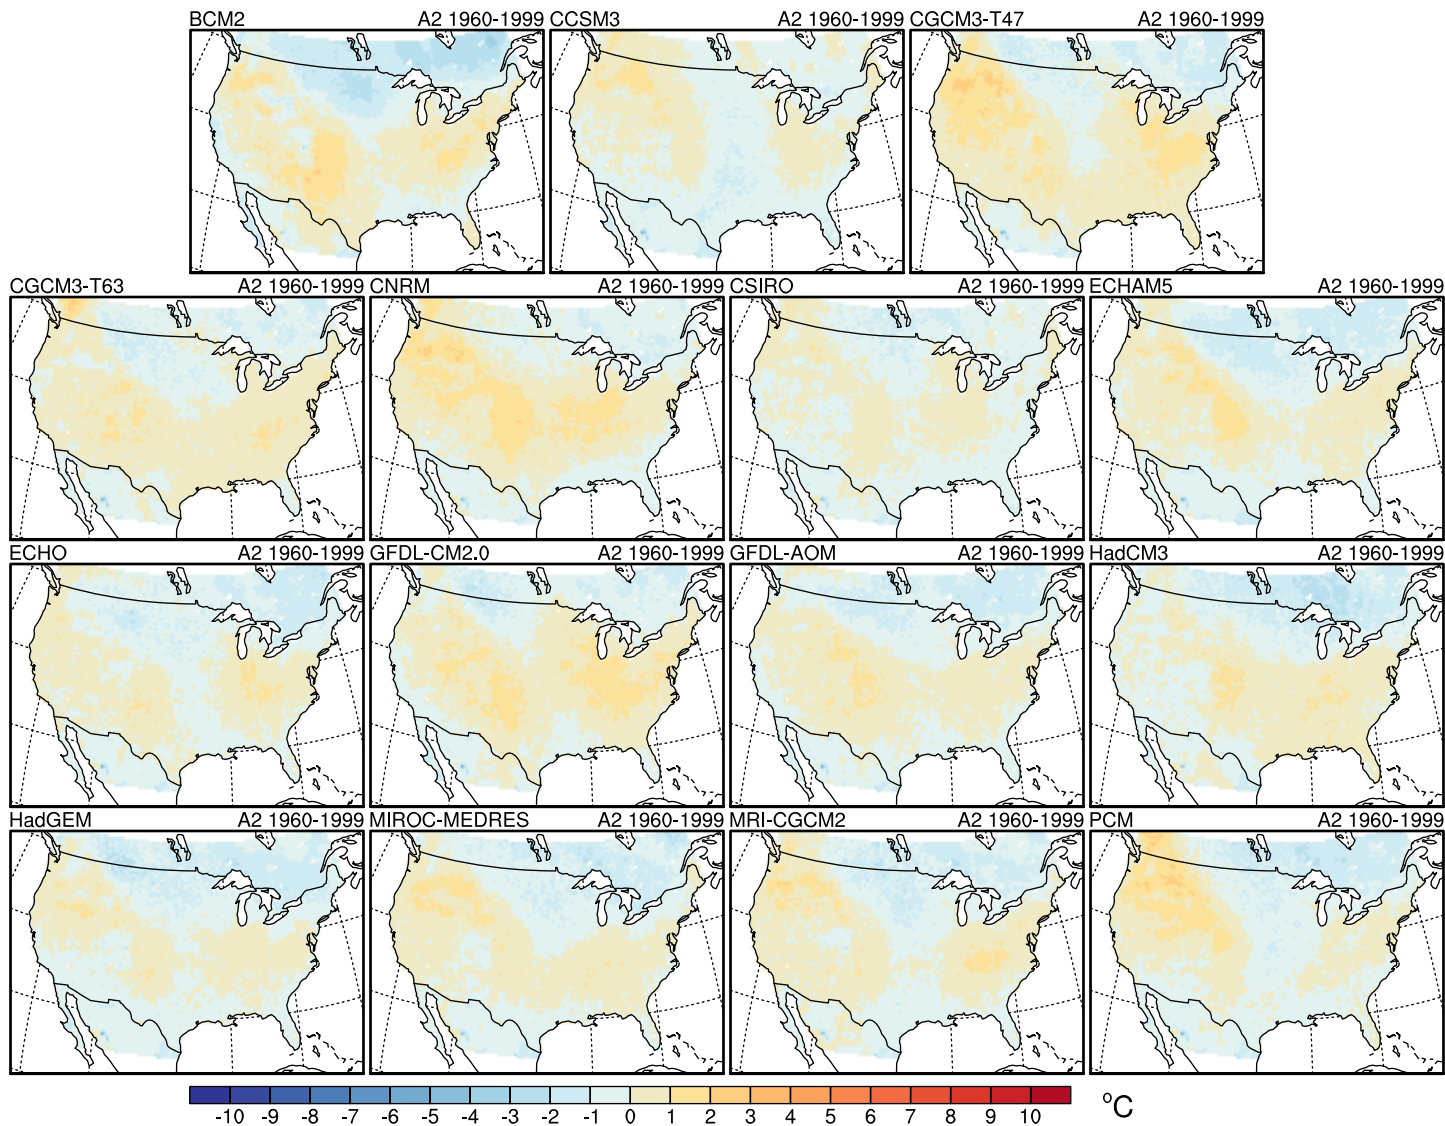
Bias in number of days with $T_{min} < 85\text{ F}$



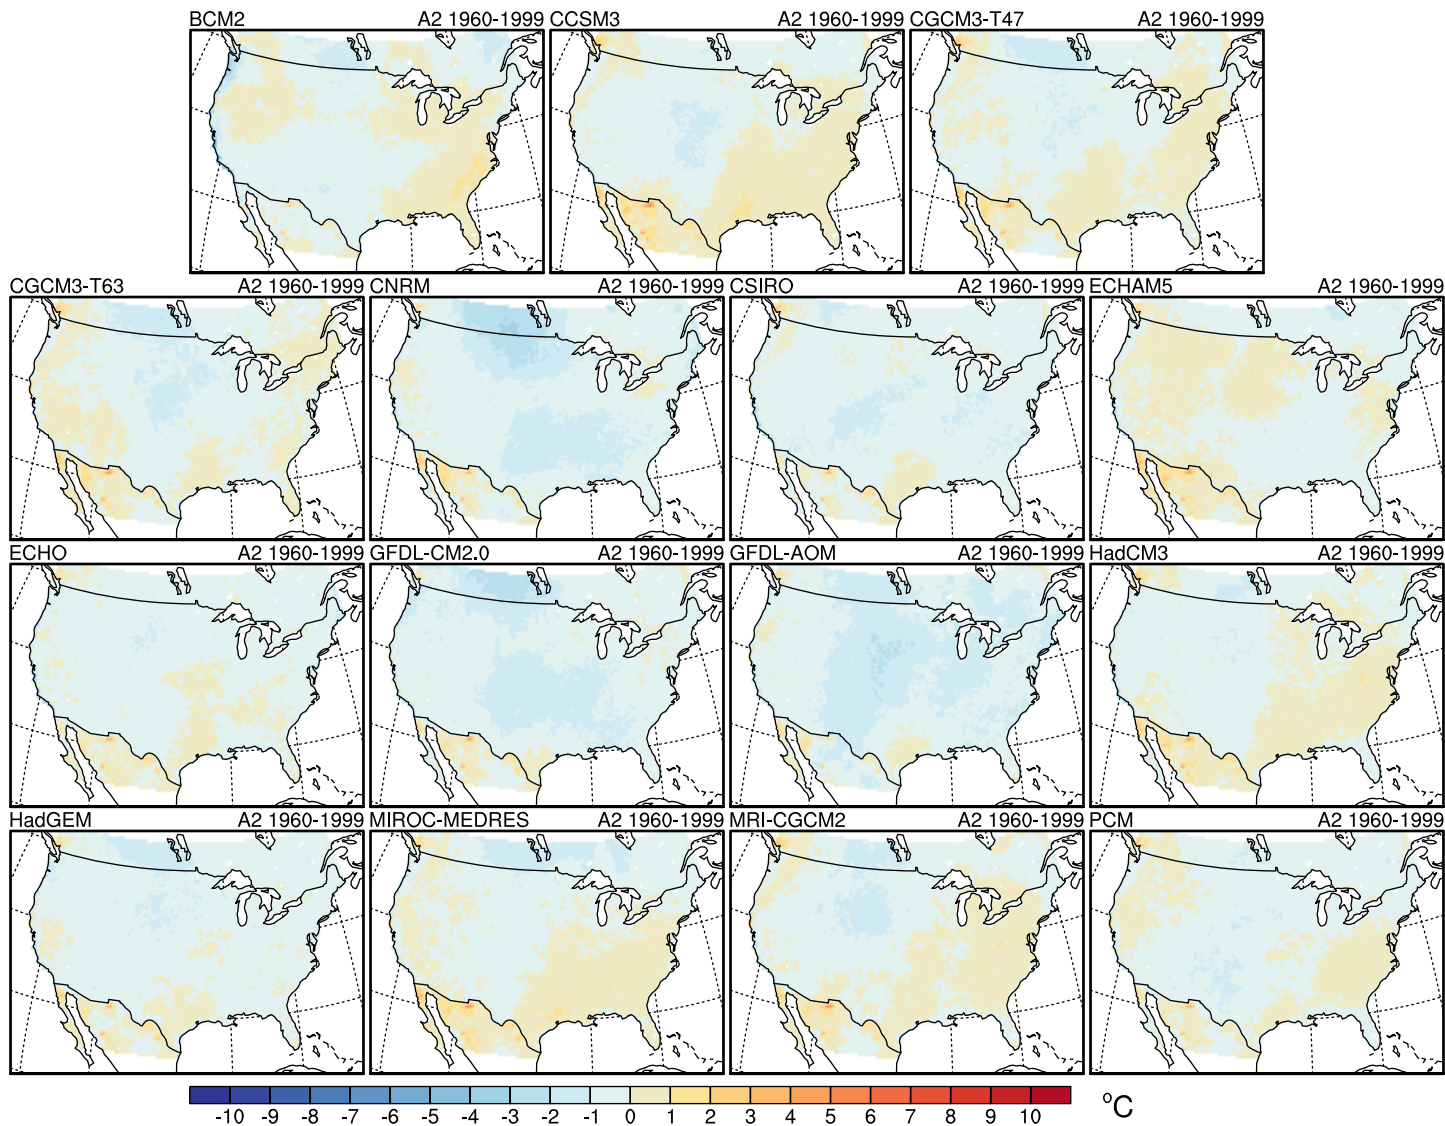
Bias in avg tmax of hottest 5-day period of the year



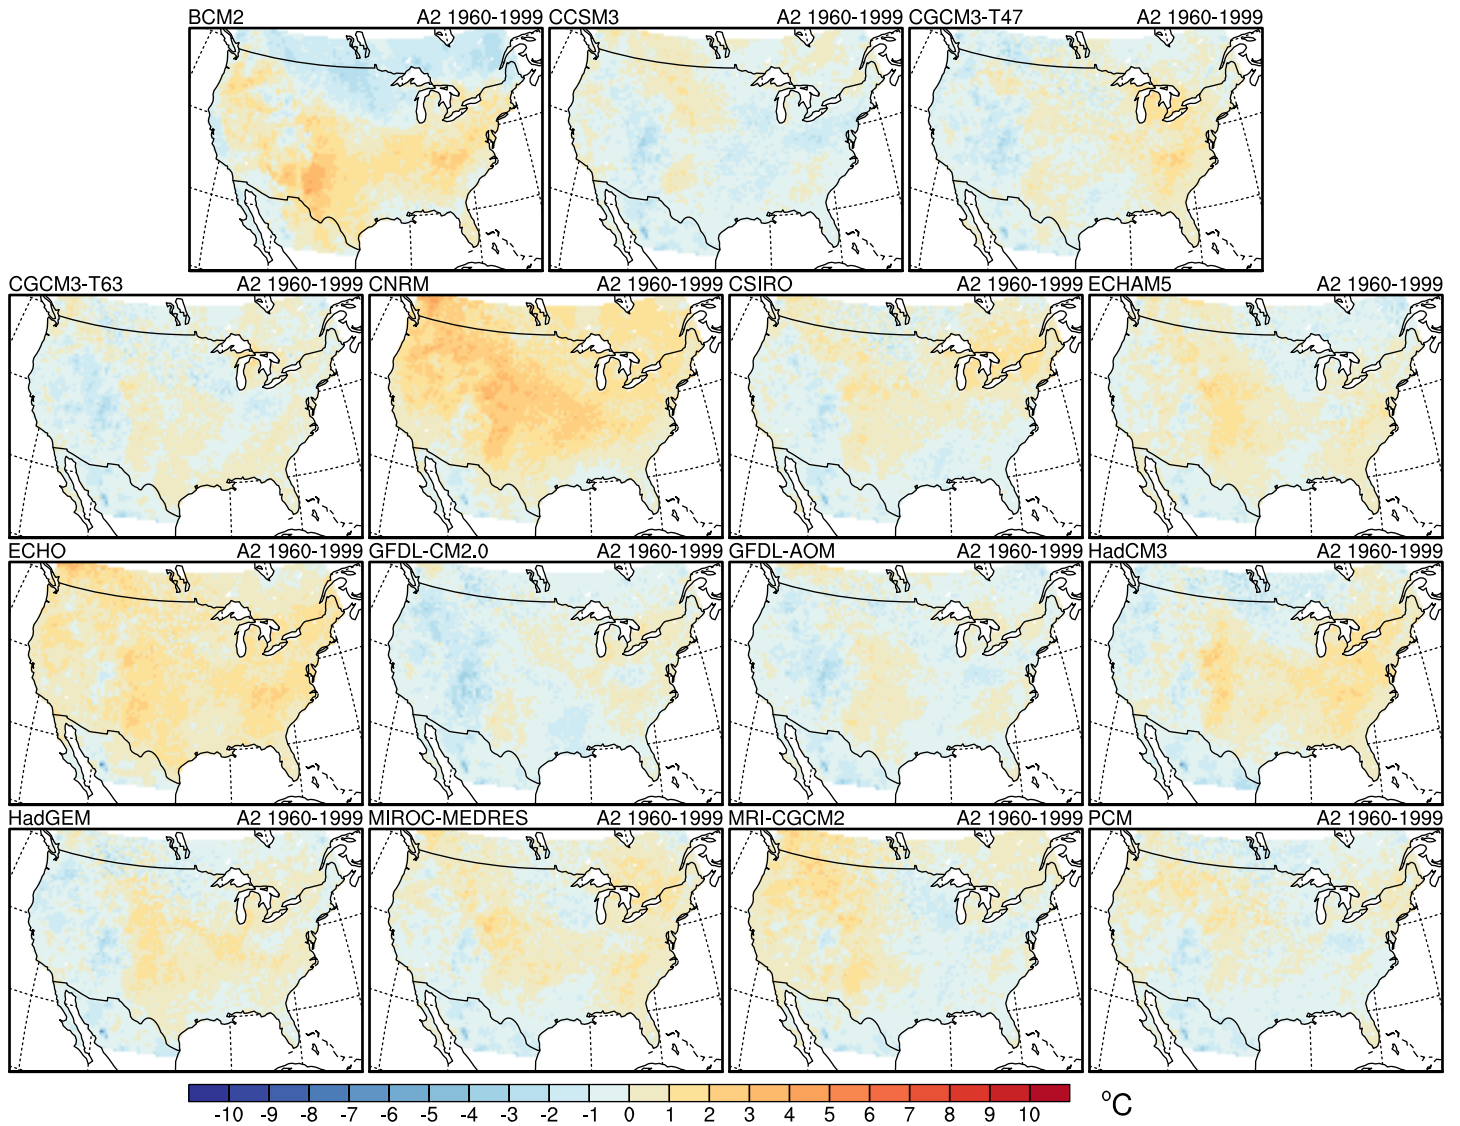
Bias in avg tmin of coldest 5-night period of the year



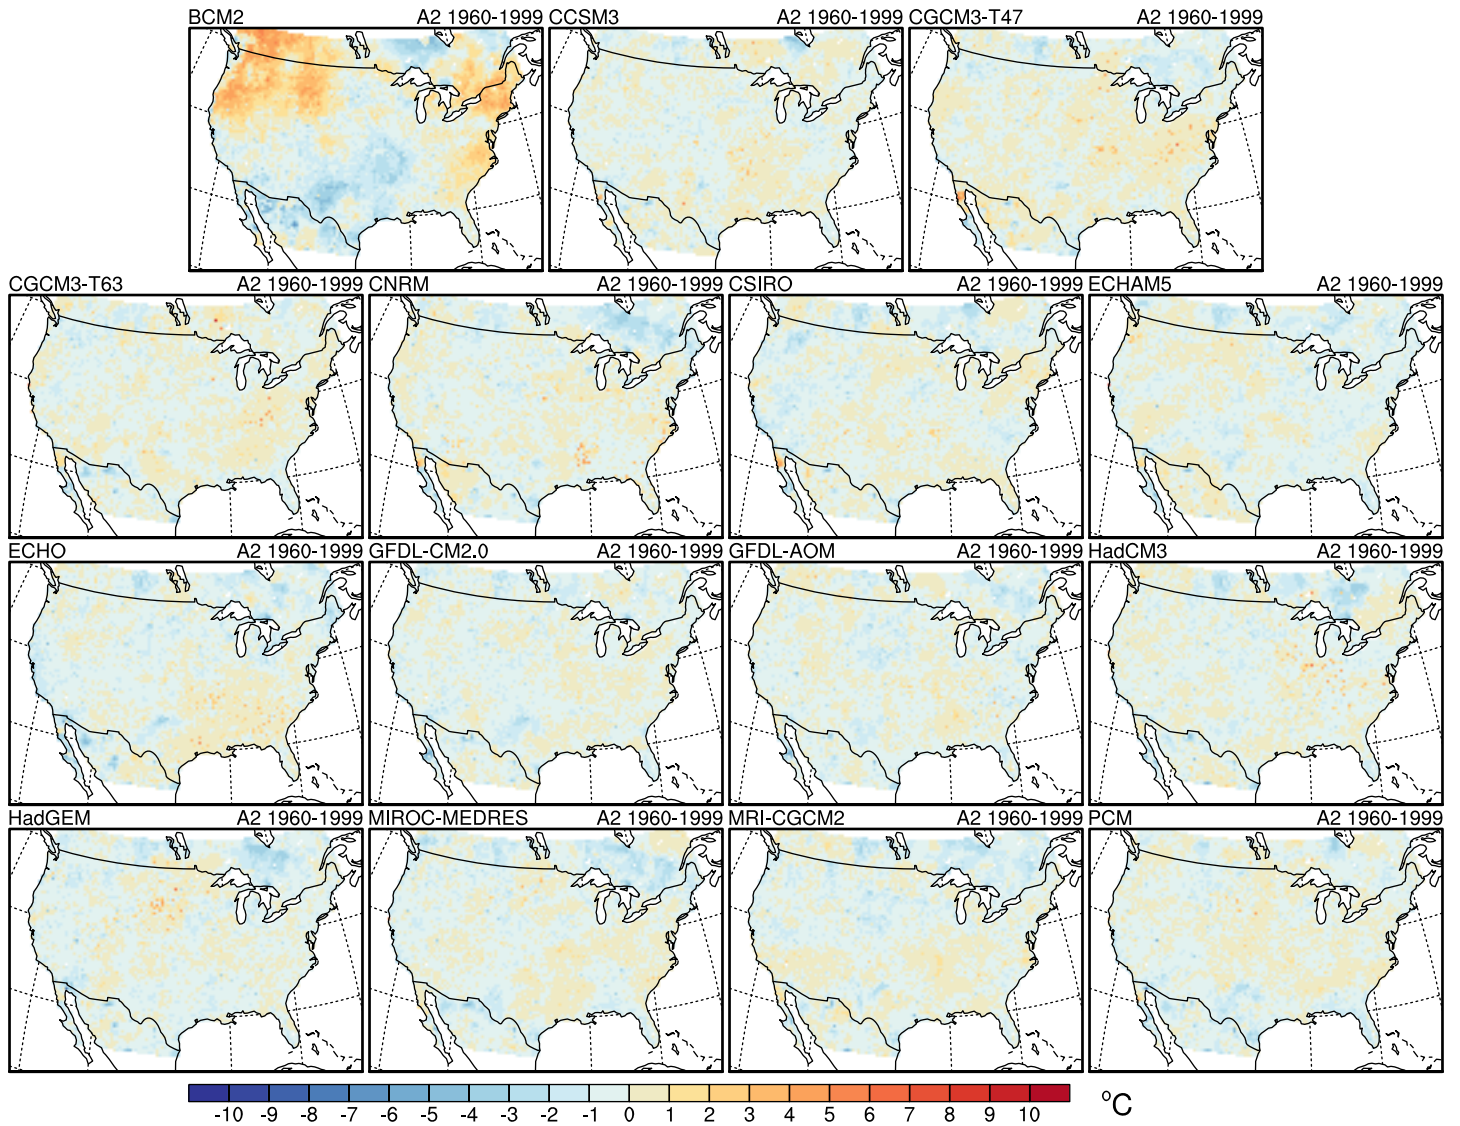
Bias in hottest day of the year



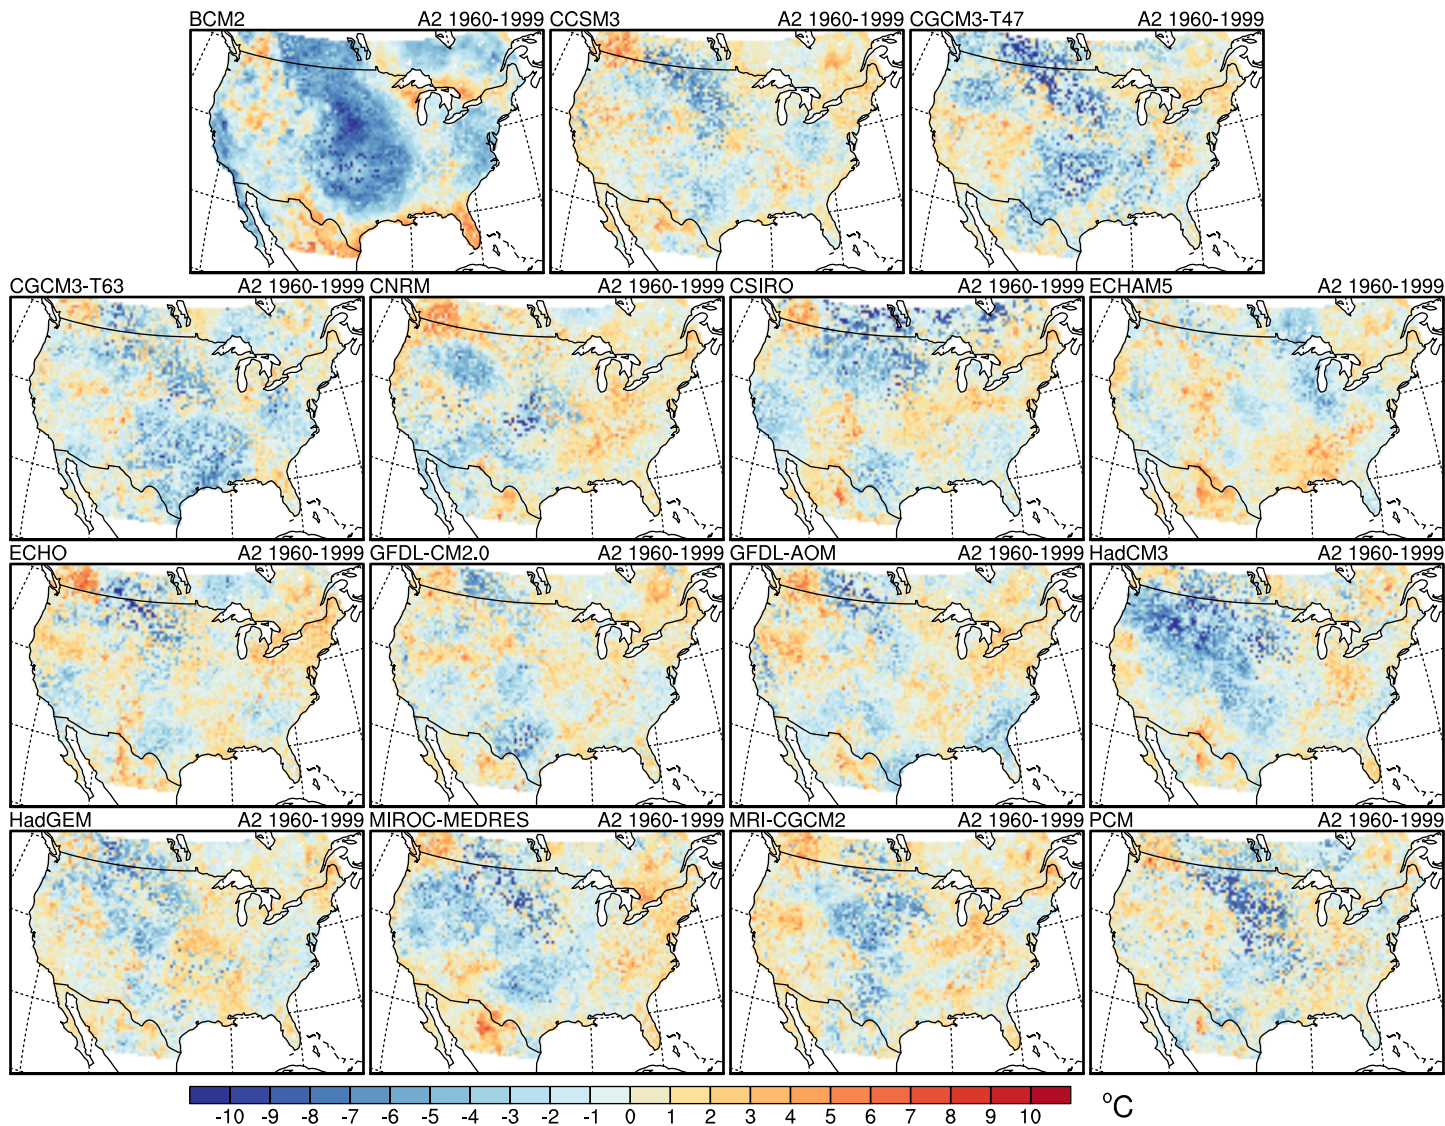
Bias in coldest night of the year



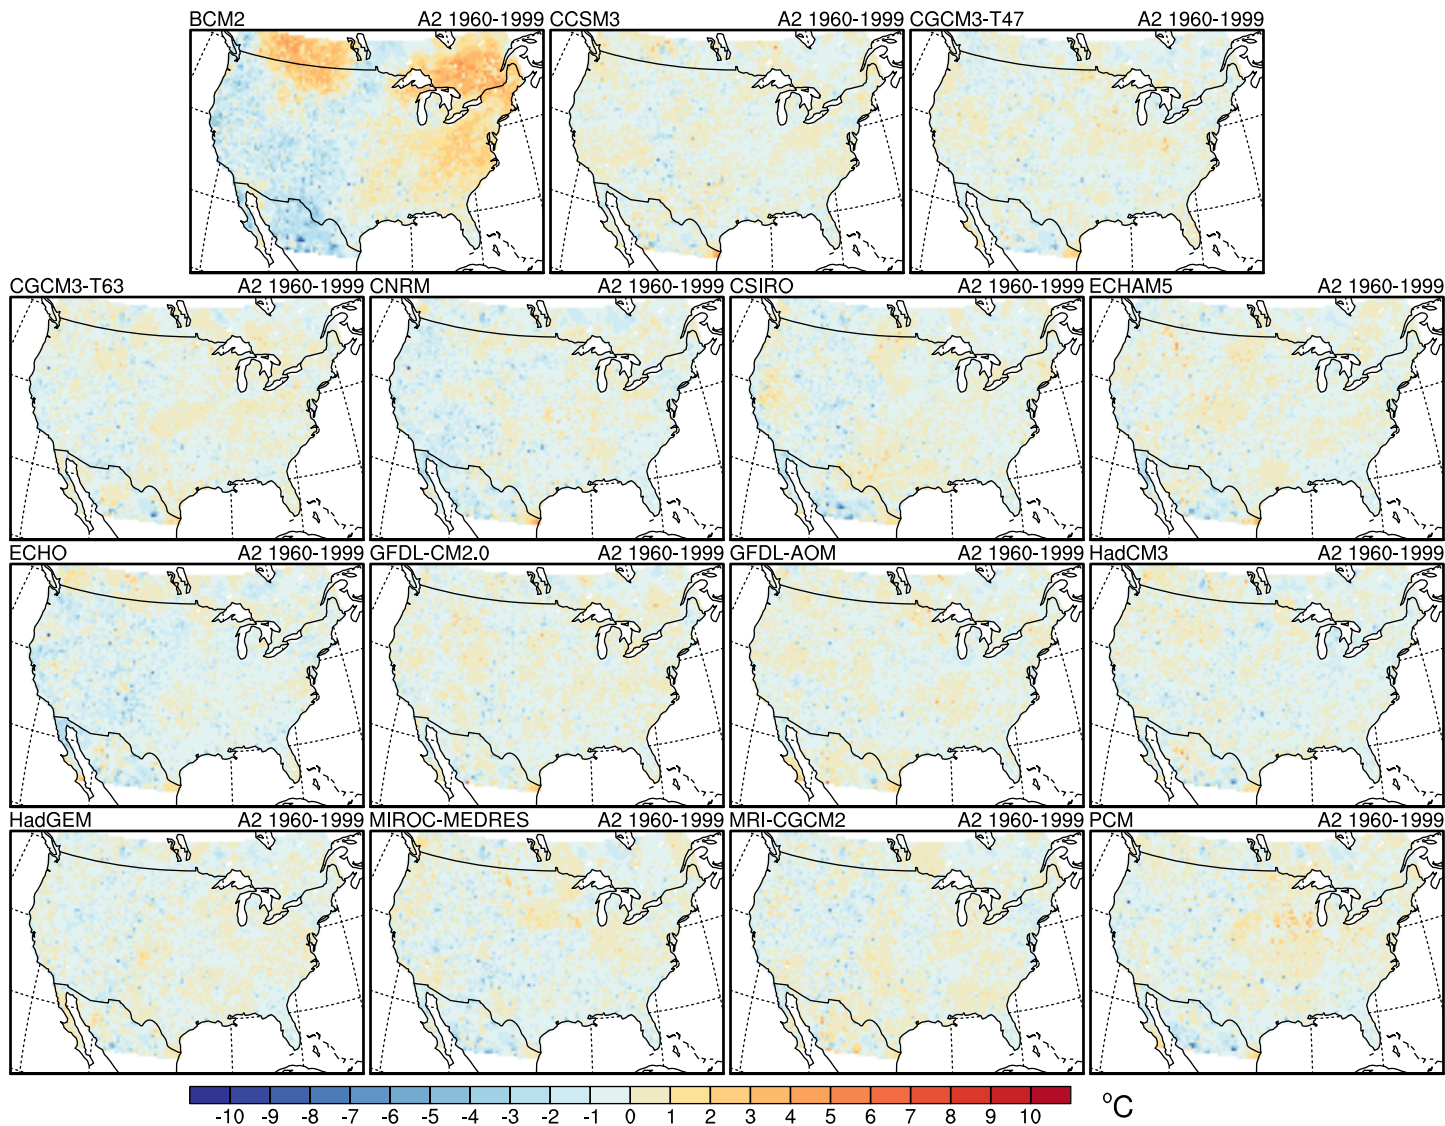
Bias in hottest day in 30 years (tmax)



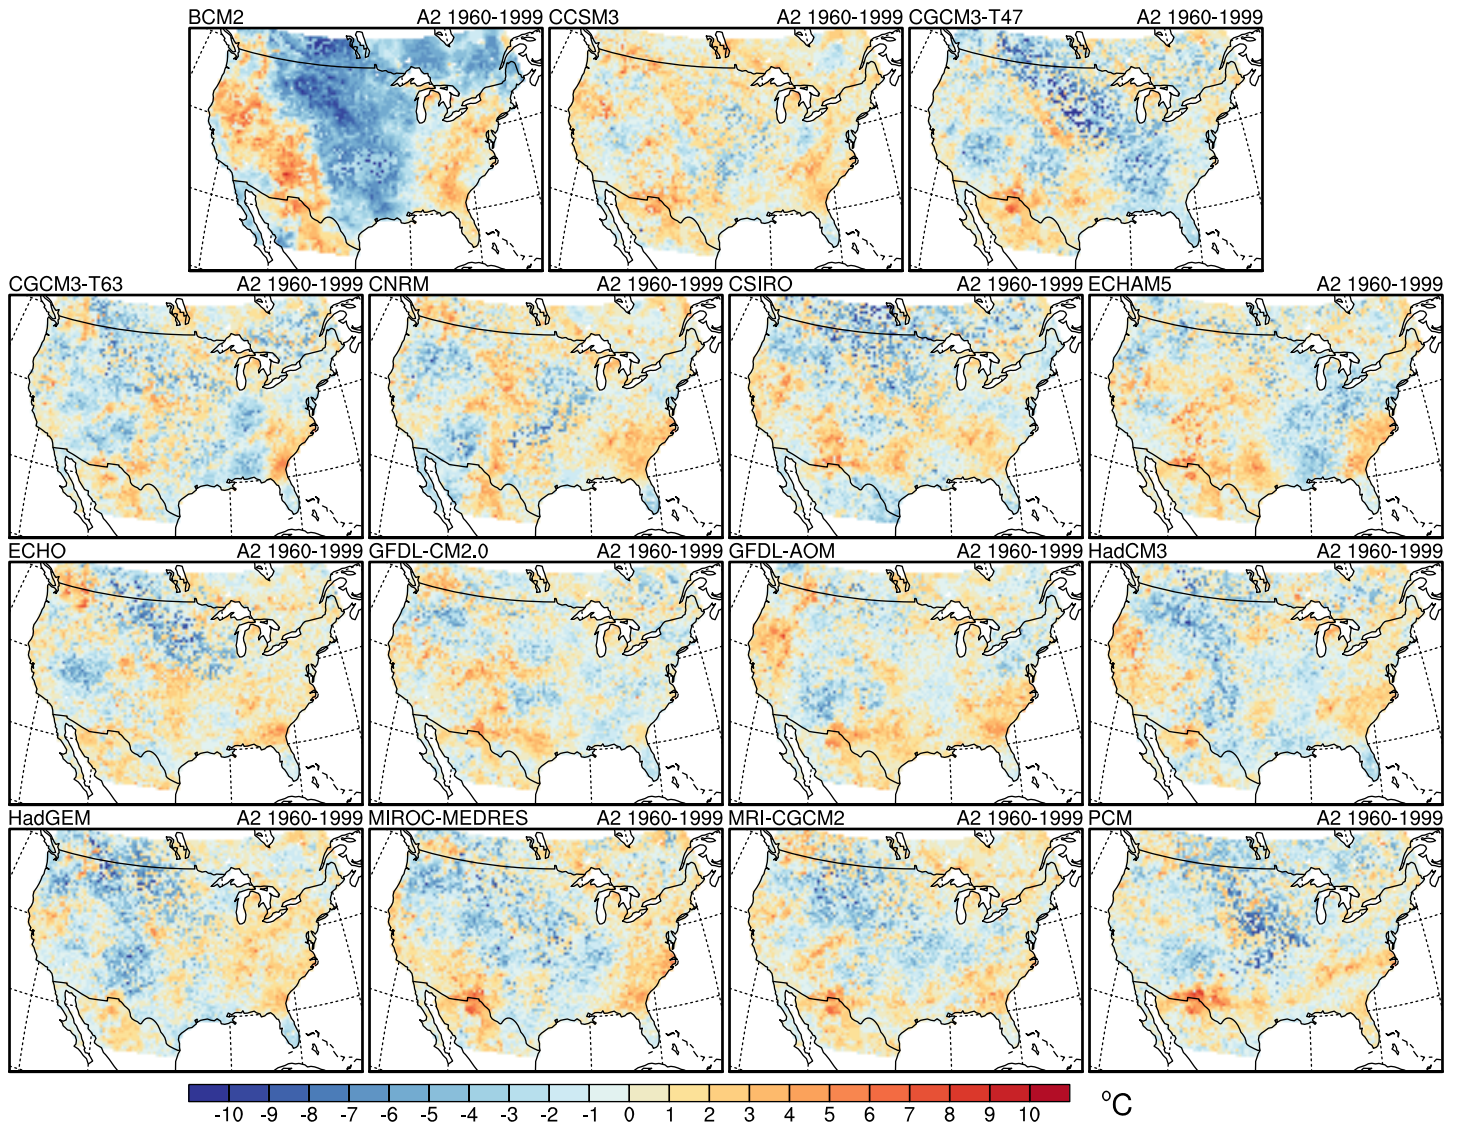
Bias in coldest day in 30 years (tmax)



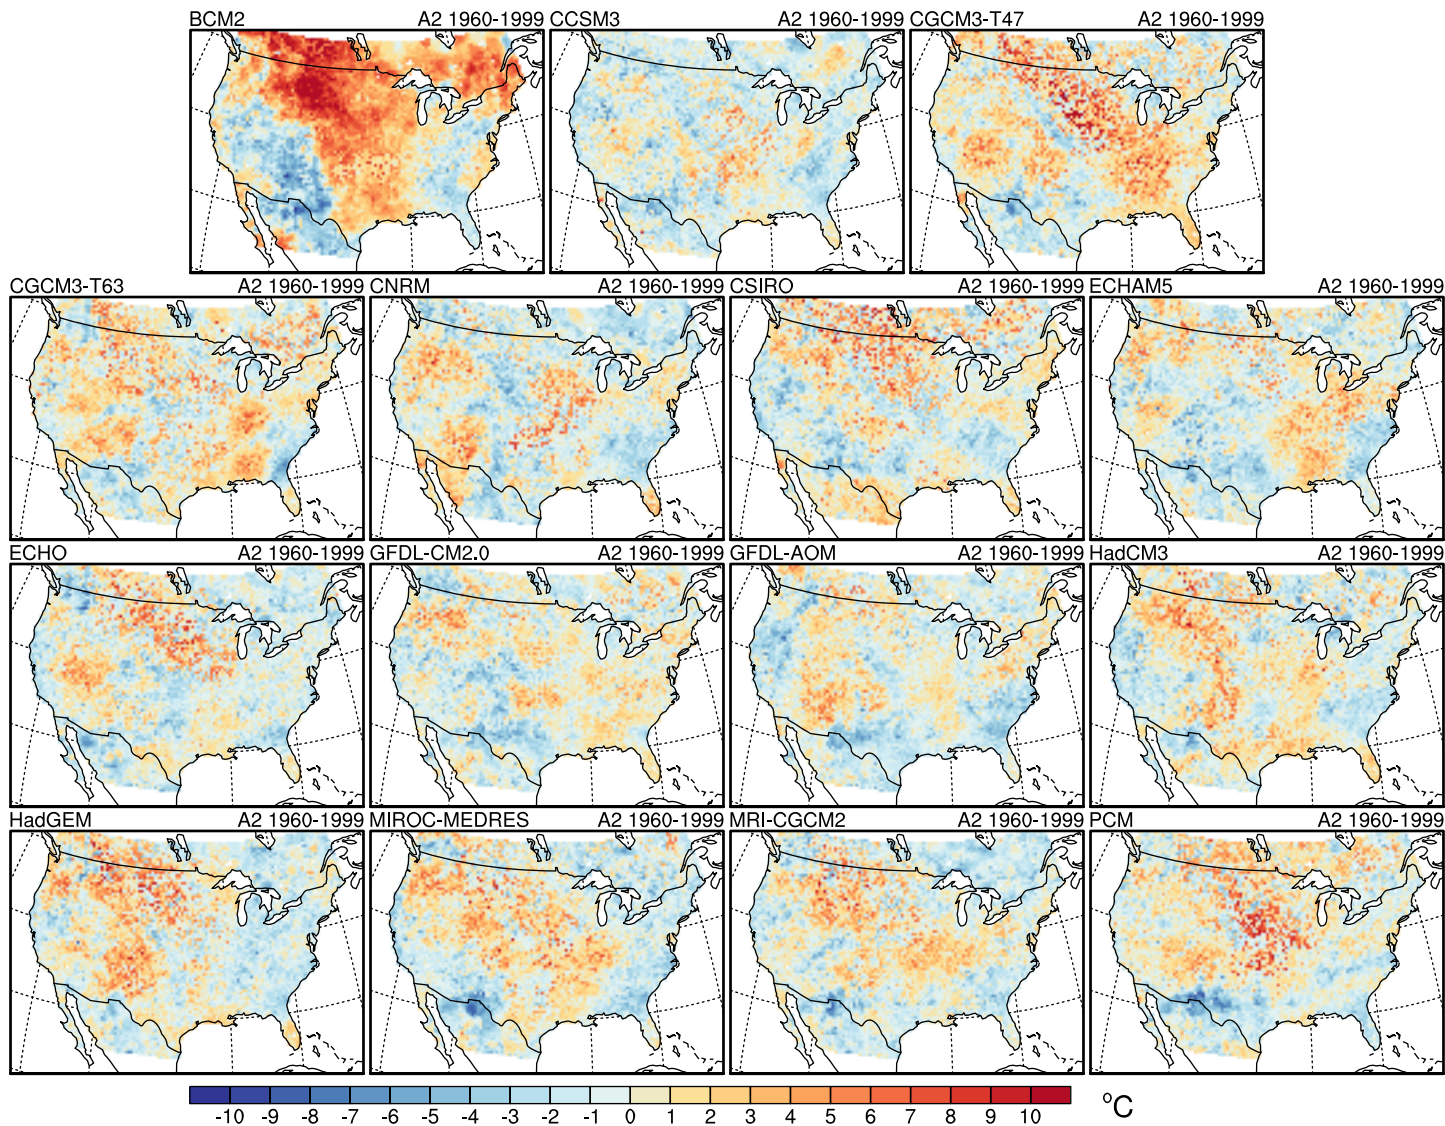
Bias in warmest night in 30 years (tmin)



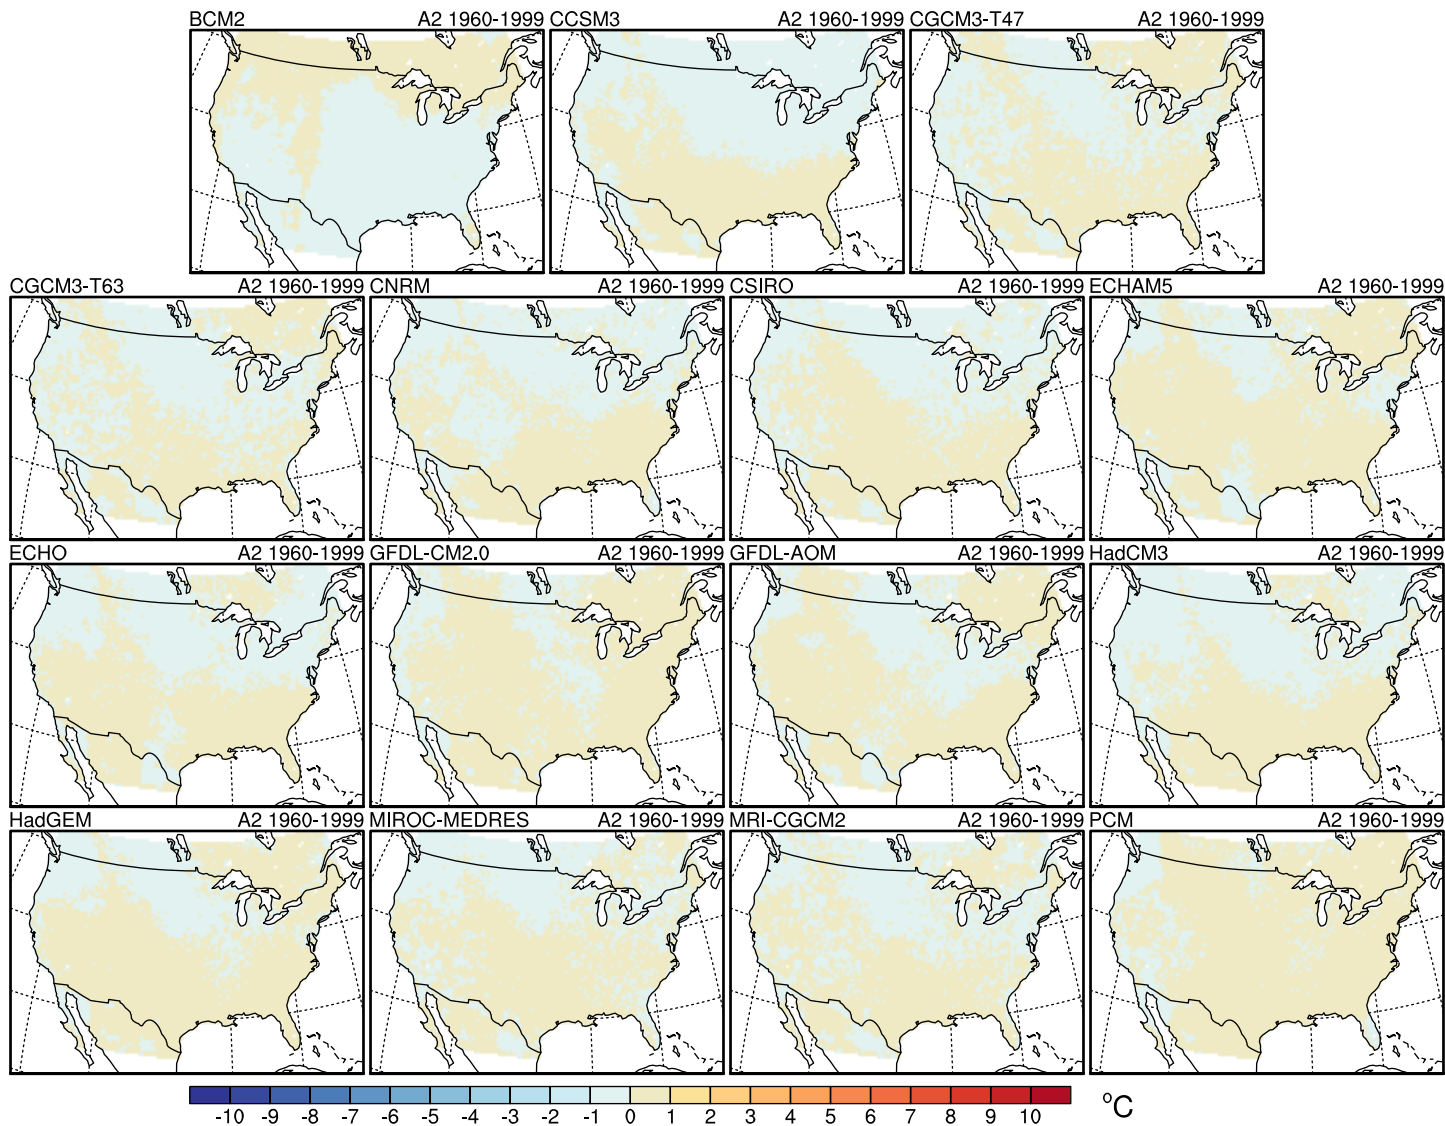
Bias in coldest night in 30 years (tmin)



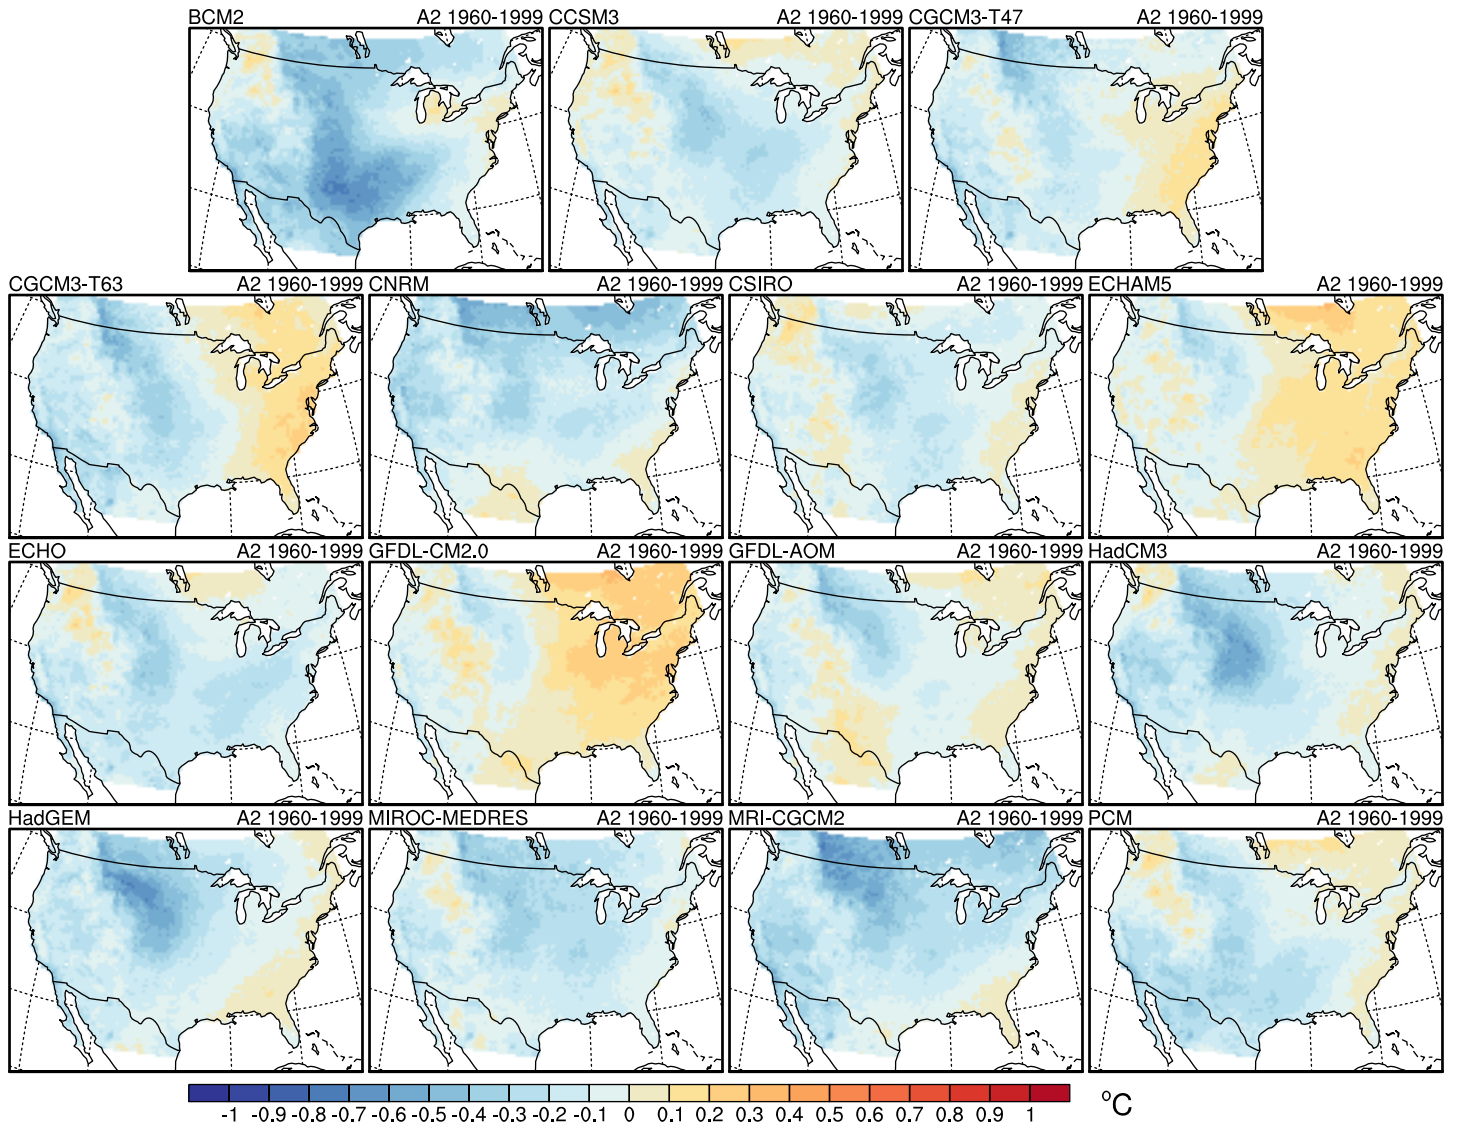
Bias in difference between hottest day and coldest night in 30 years



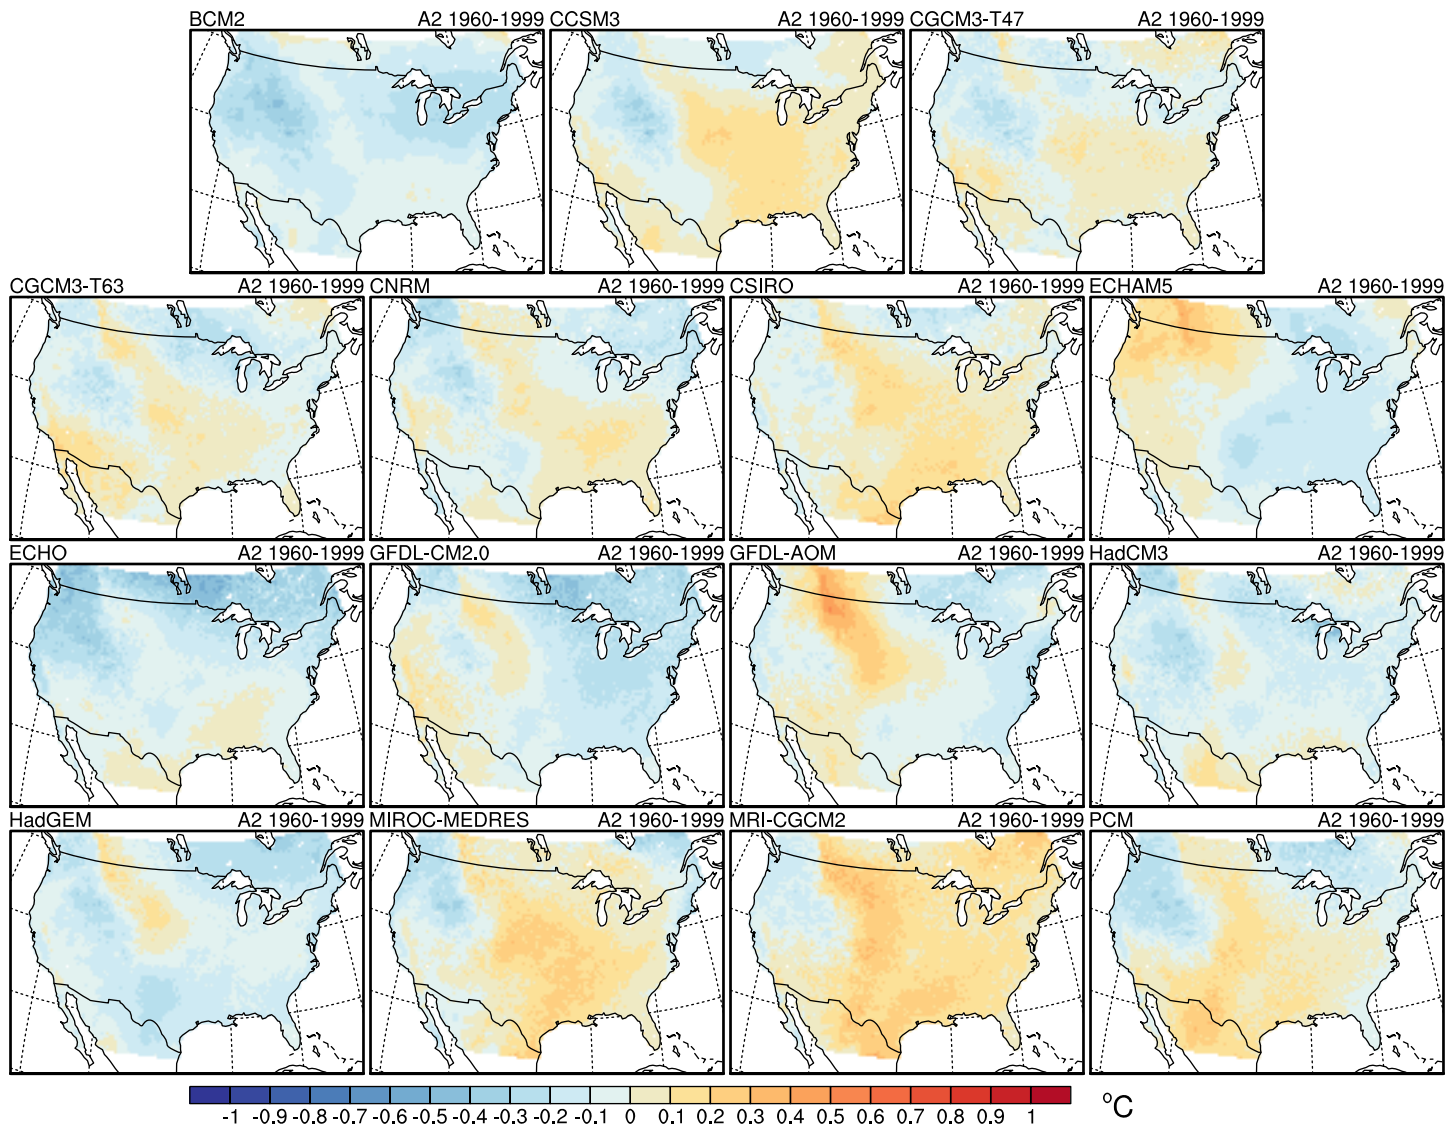
Bias in difference between mean tmax and tmin over 30 years




Bias in mean winter (DJF) Tmax



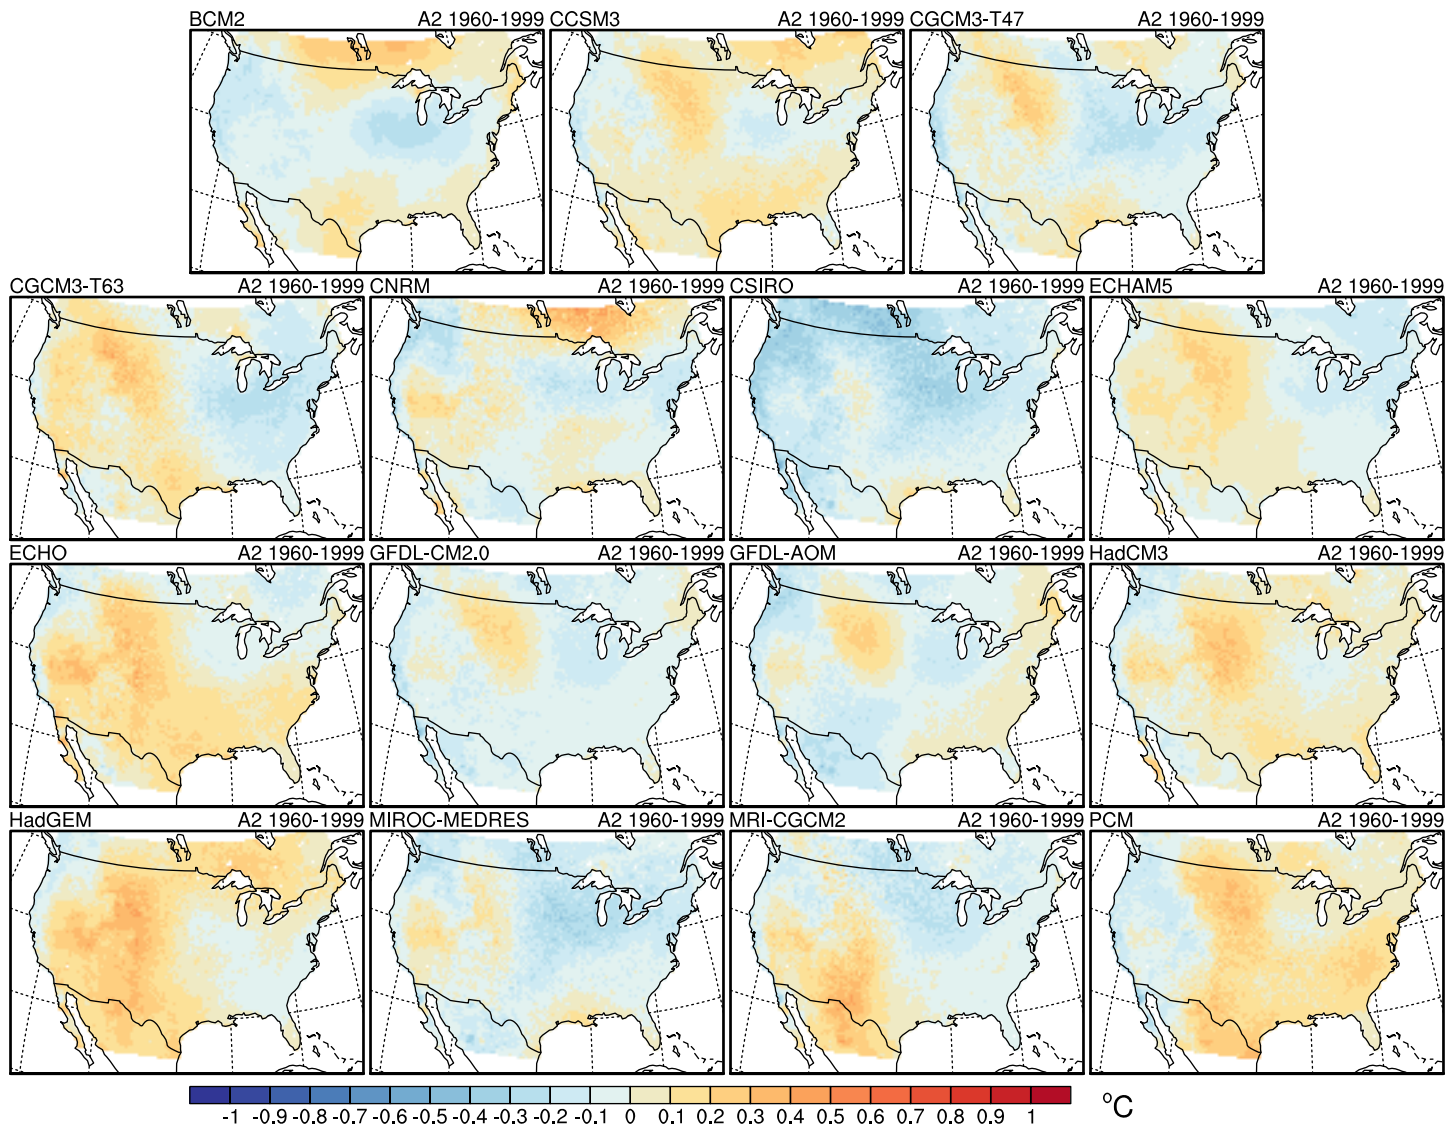
Bias in mean spring (MAM) Tmax




Bias in mean summer (JJA) Tmax



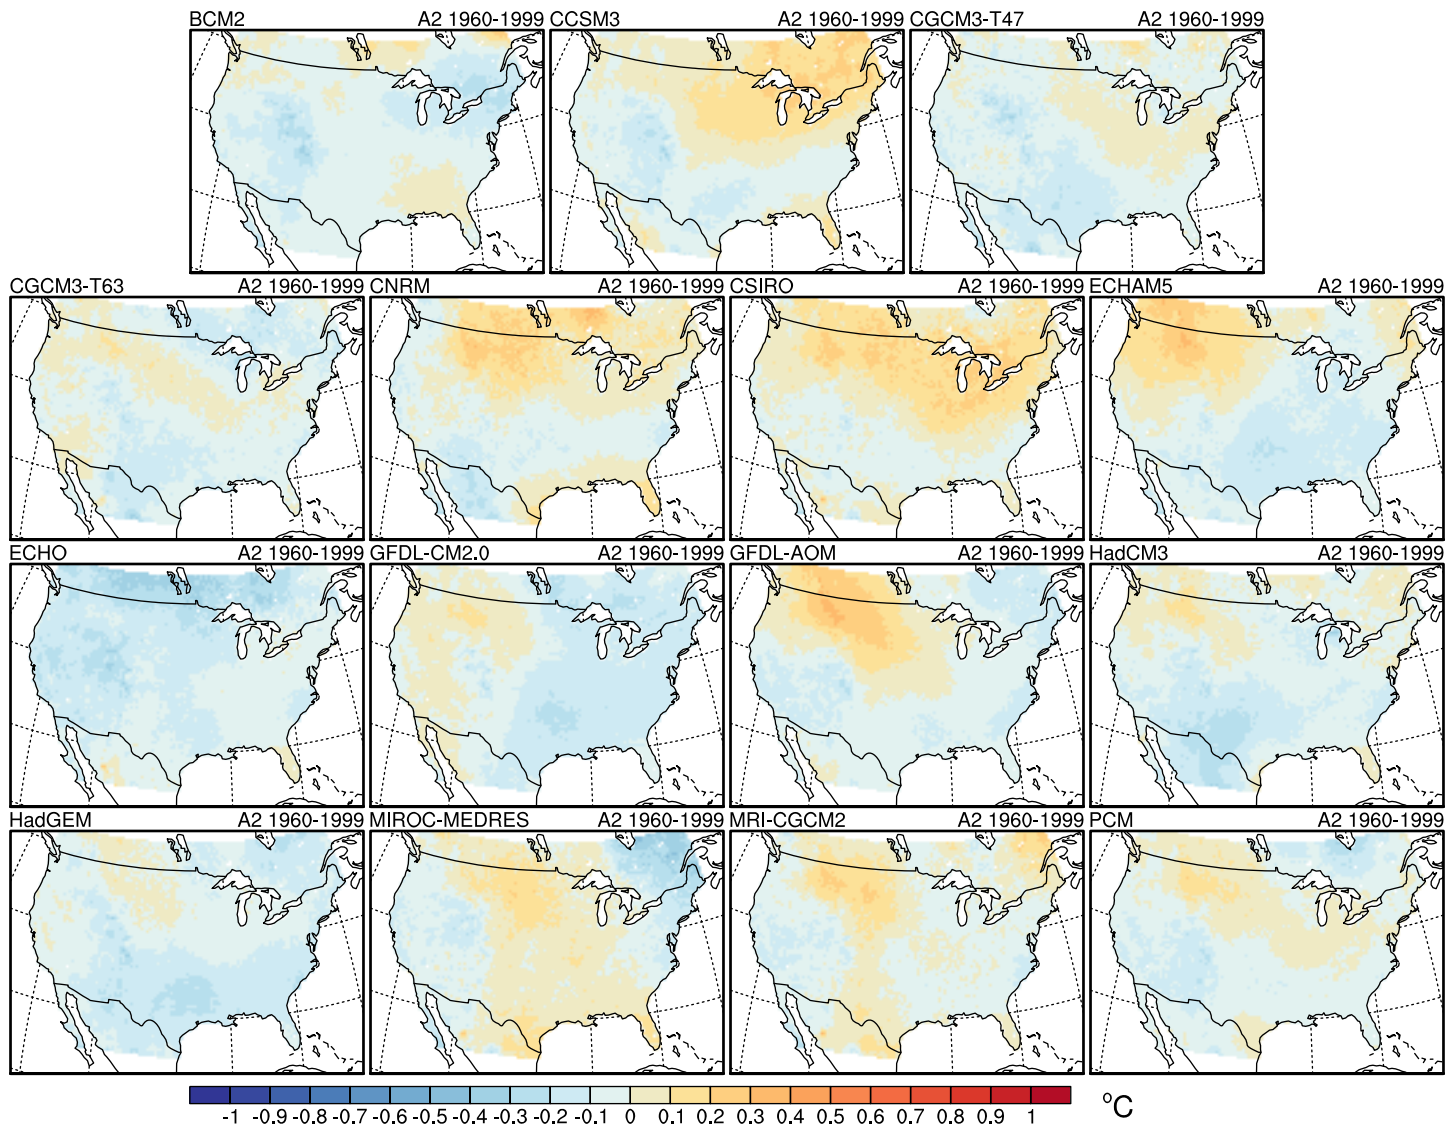
Bias in mean fall (SON) Tmax



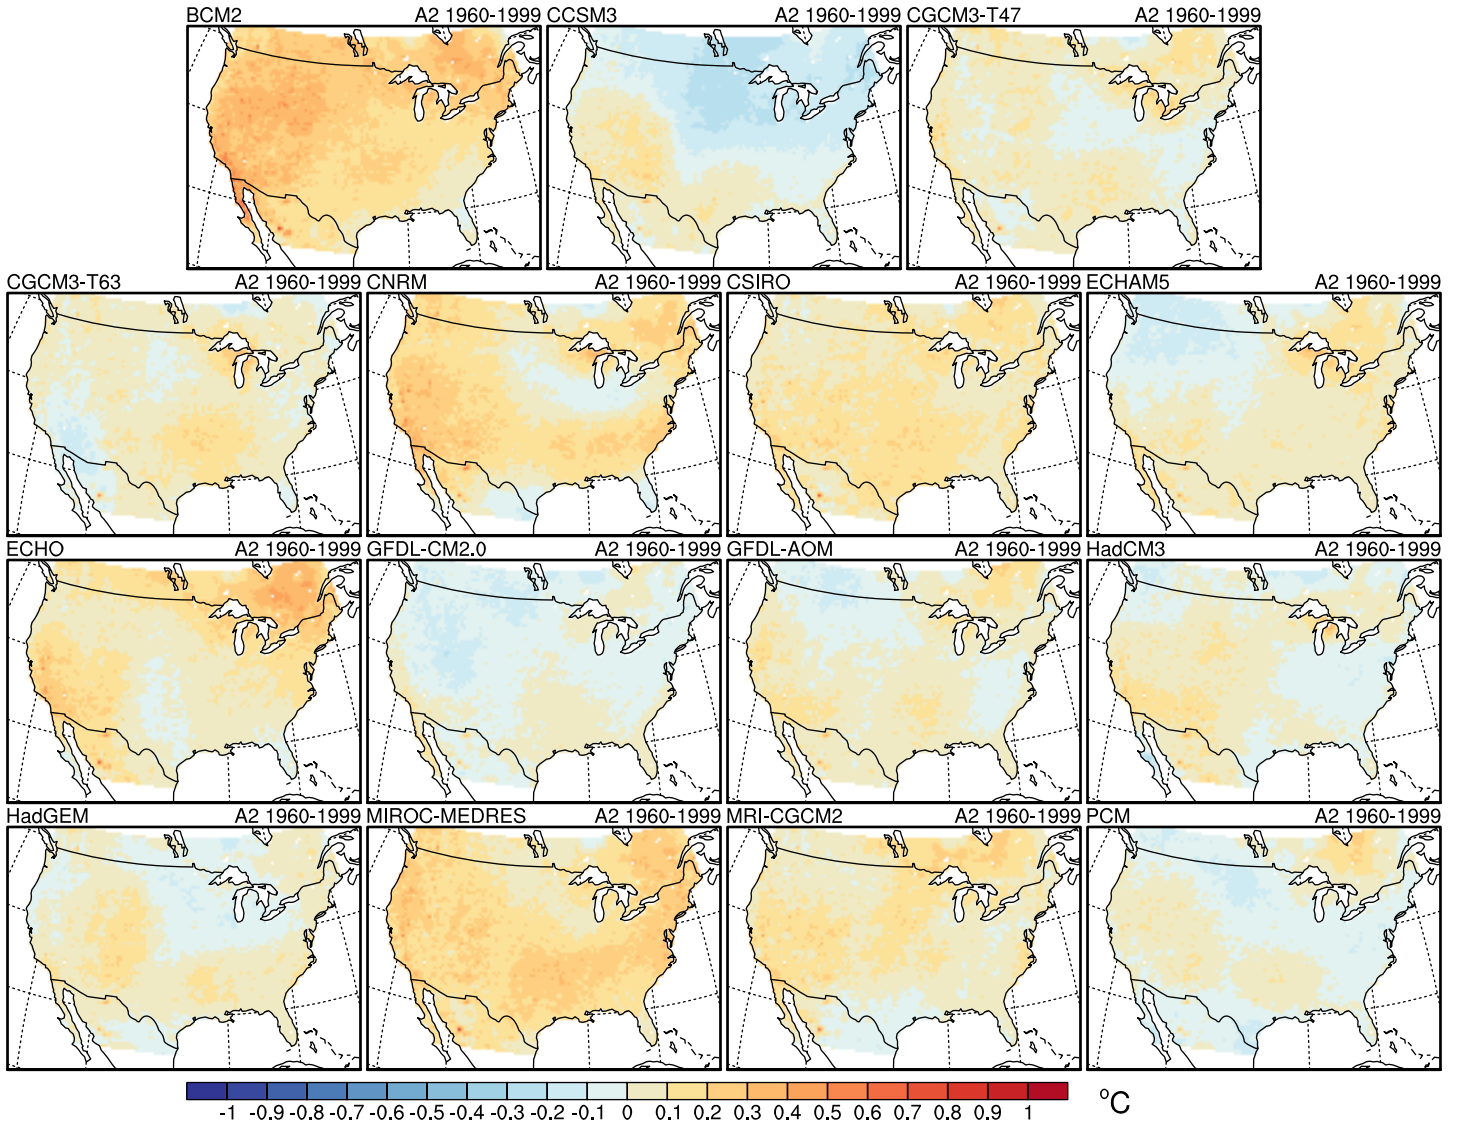
Bias in mean winter (DJF) Tmin



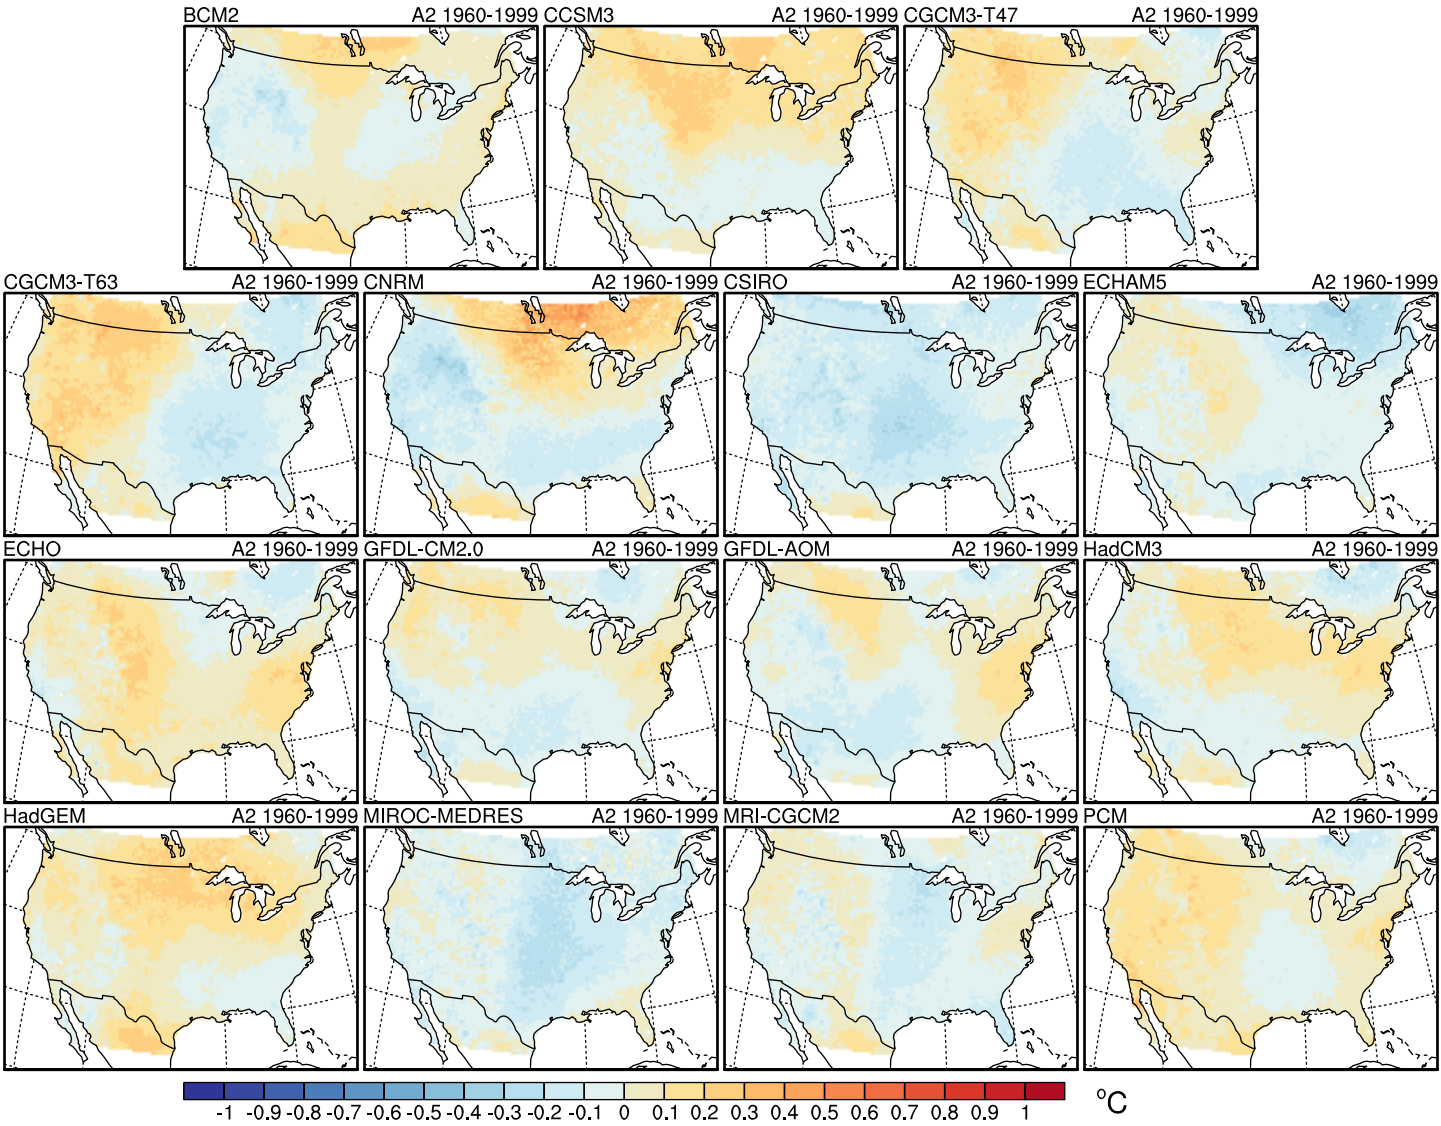
Bias in mean spring (MAM) Tmin




Bias in mean summer (JJA) Tmin



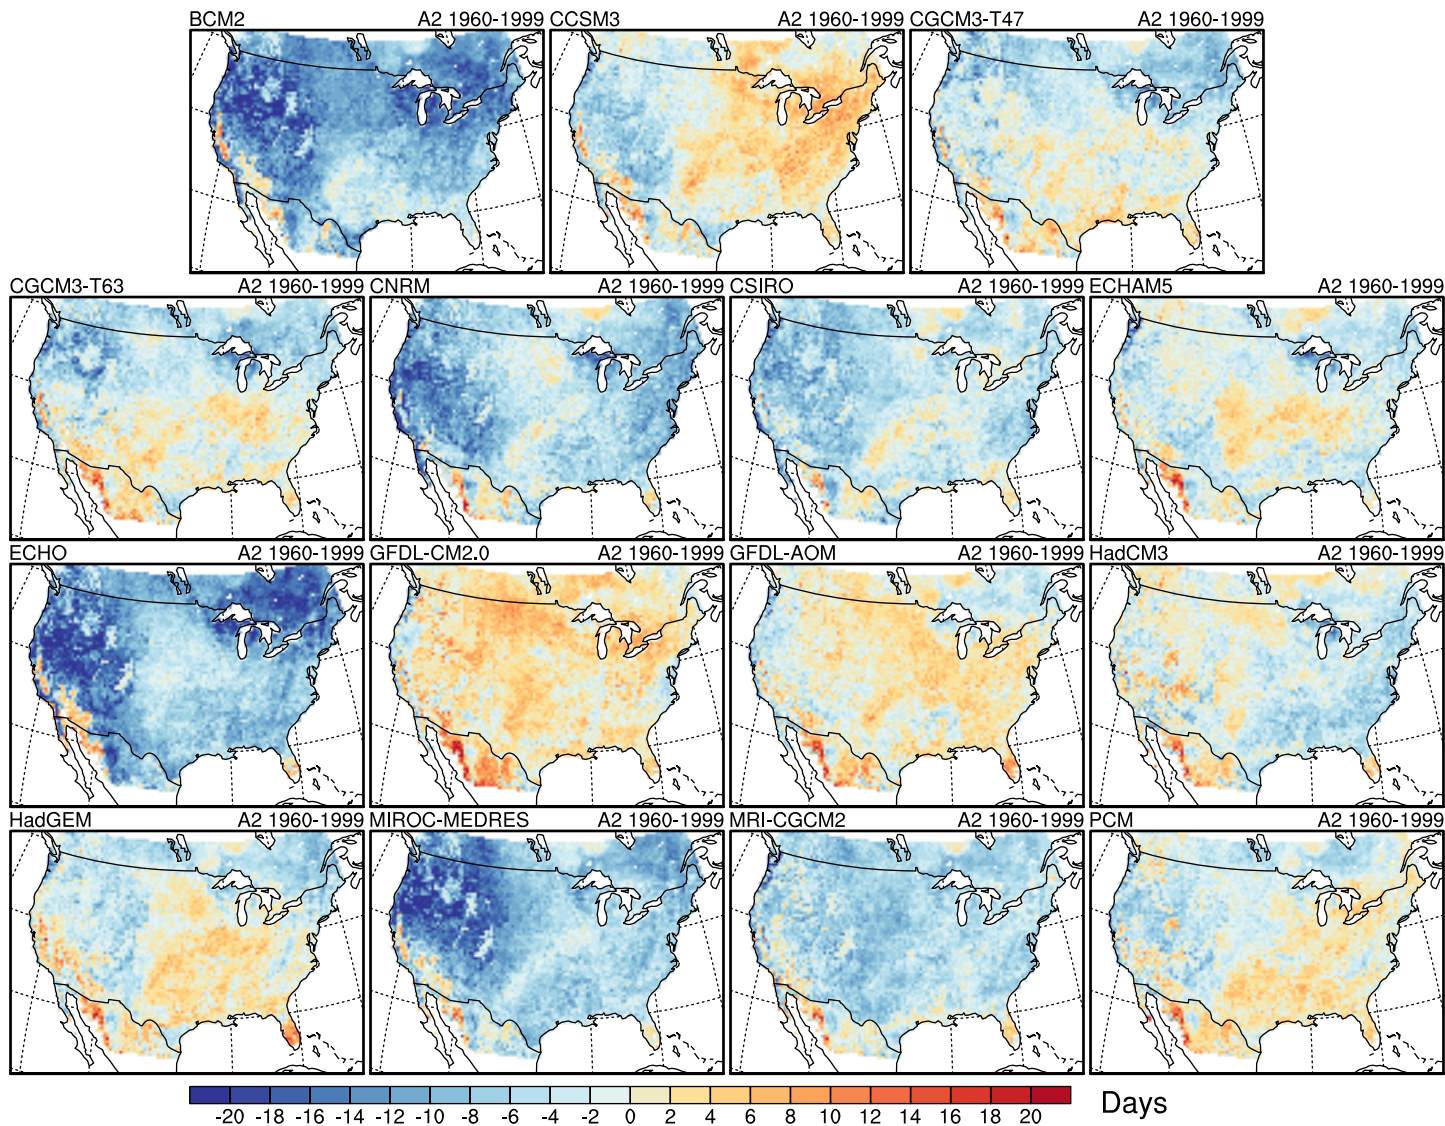
Bias in mean fall (SON) Tmin



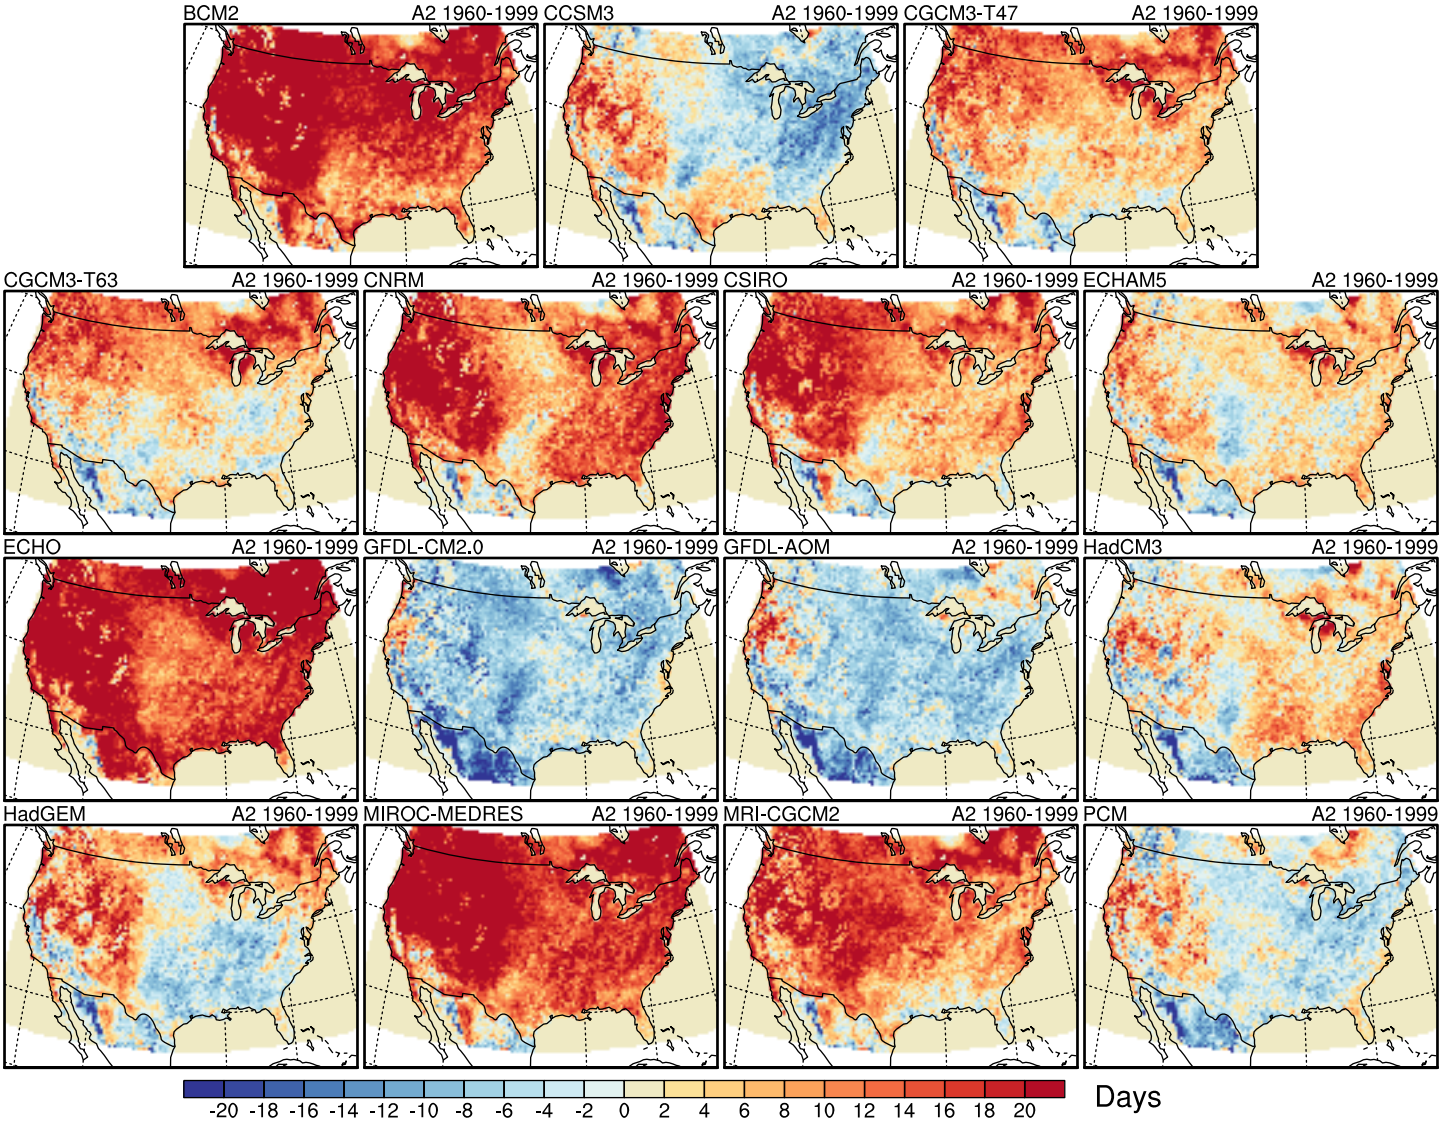
Bias in annual cumulative degree-days (threshold of 65 F)



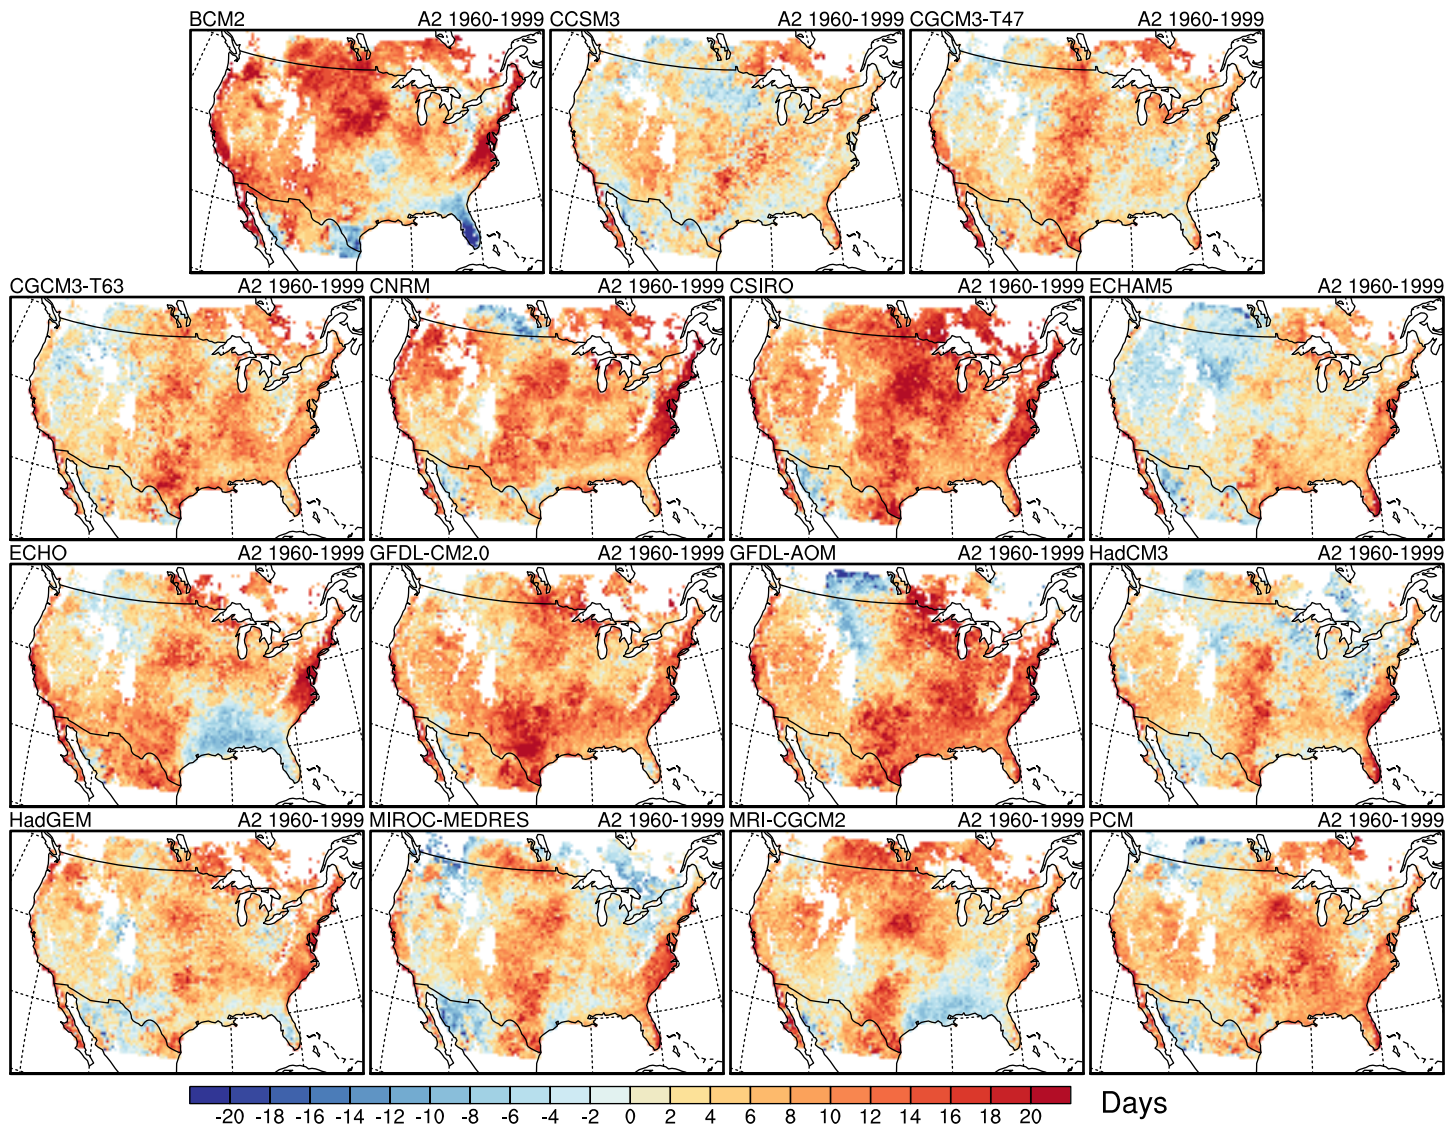
Bias in onset of spring (last day with tmin < 32F)



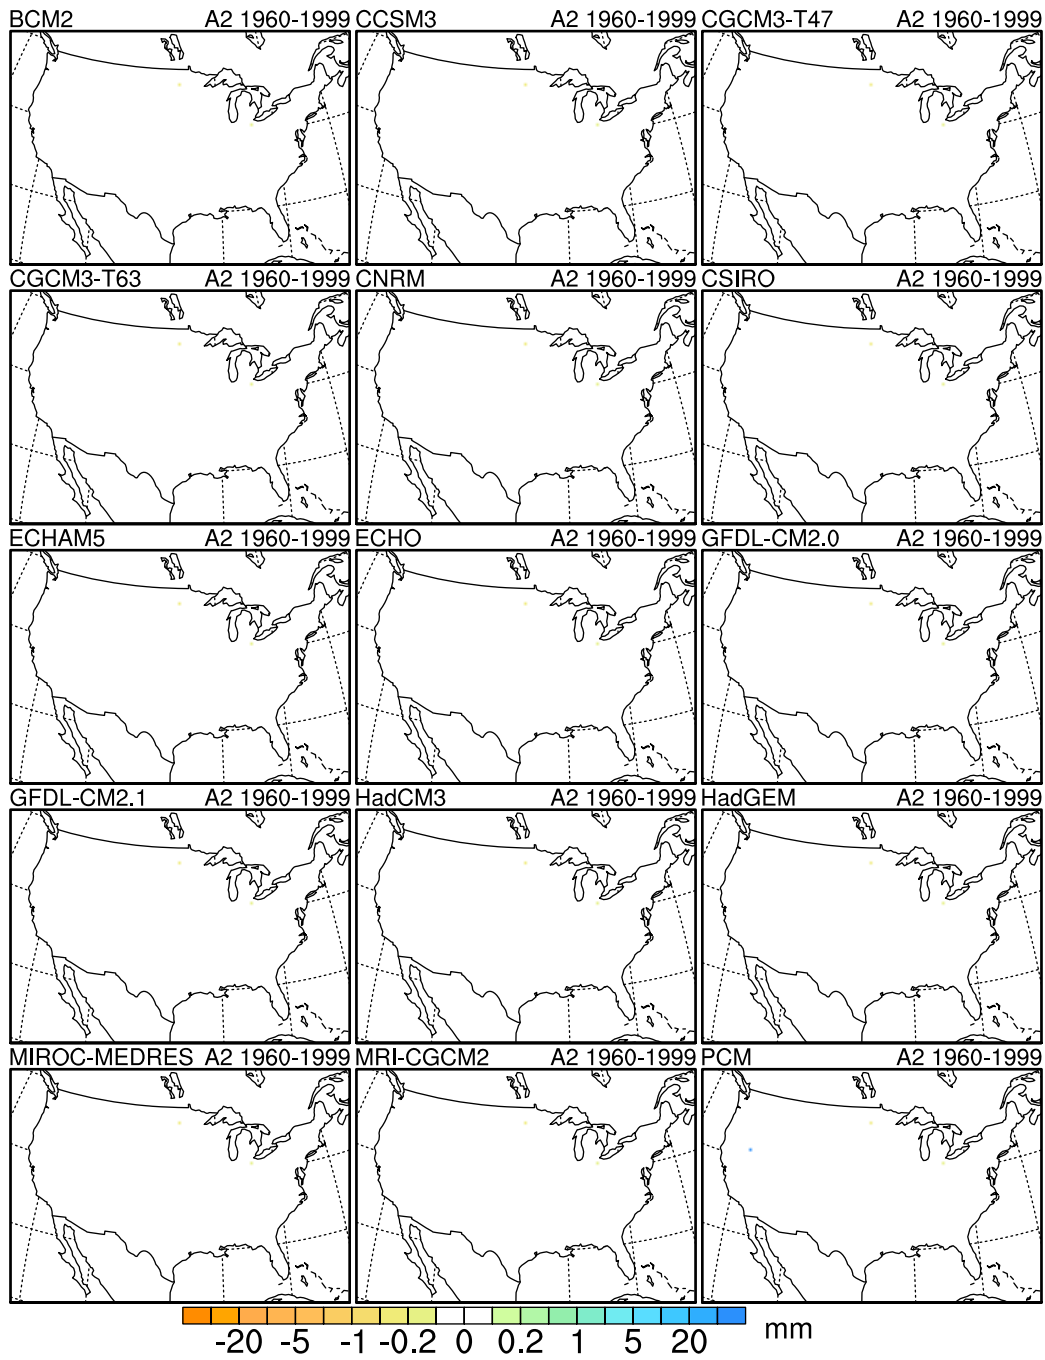
Bias in growing season length (last spring frost to first fall frost)




Bias in onset of summer (first day with tmax > 90F)



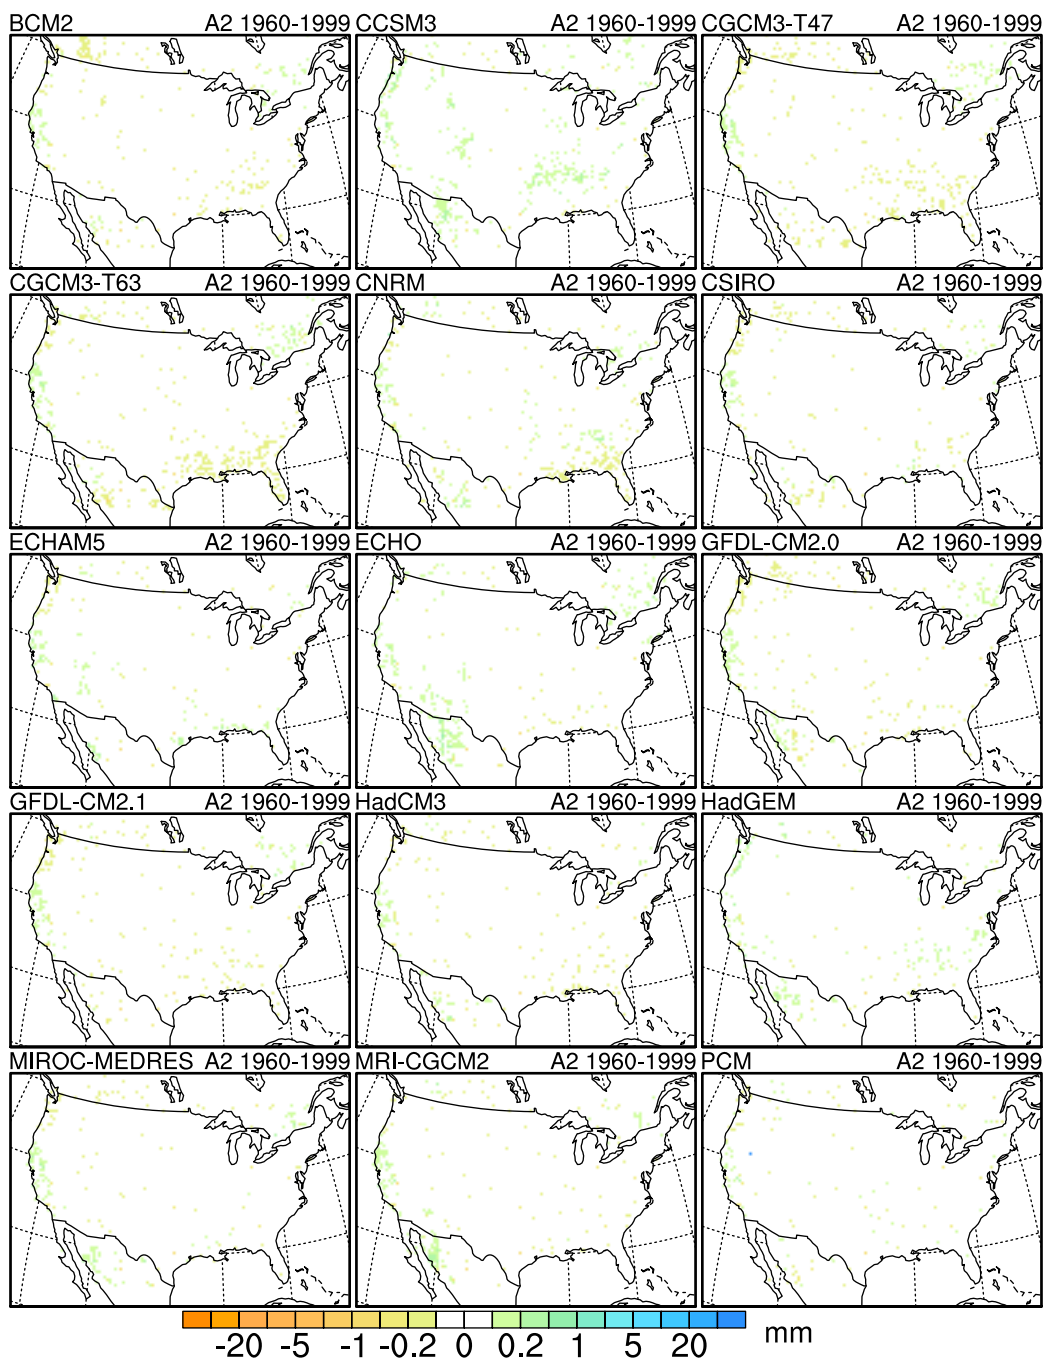
Pr Bias in 0.1th Quantile



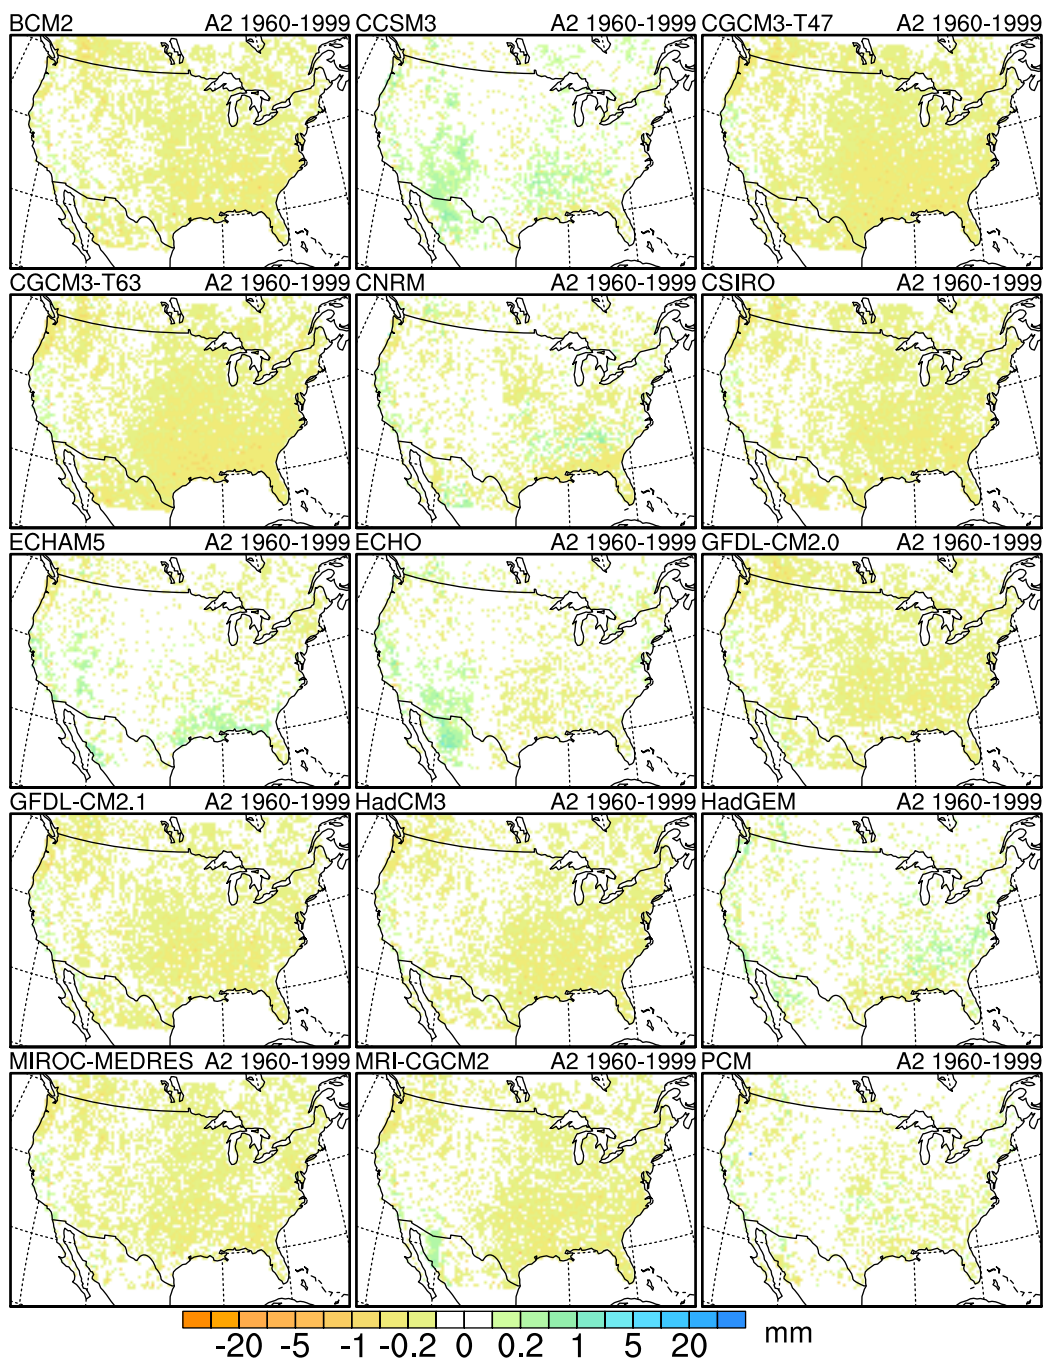
Pr Bias in 1st Quantile



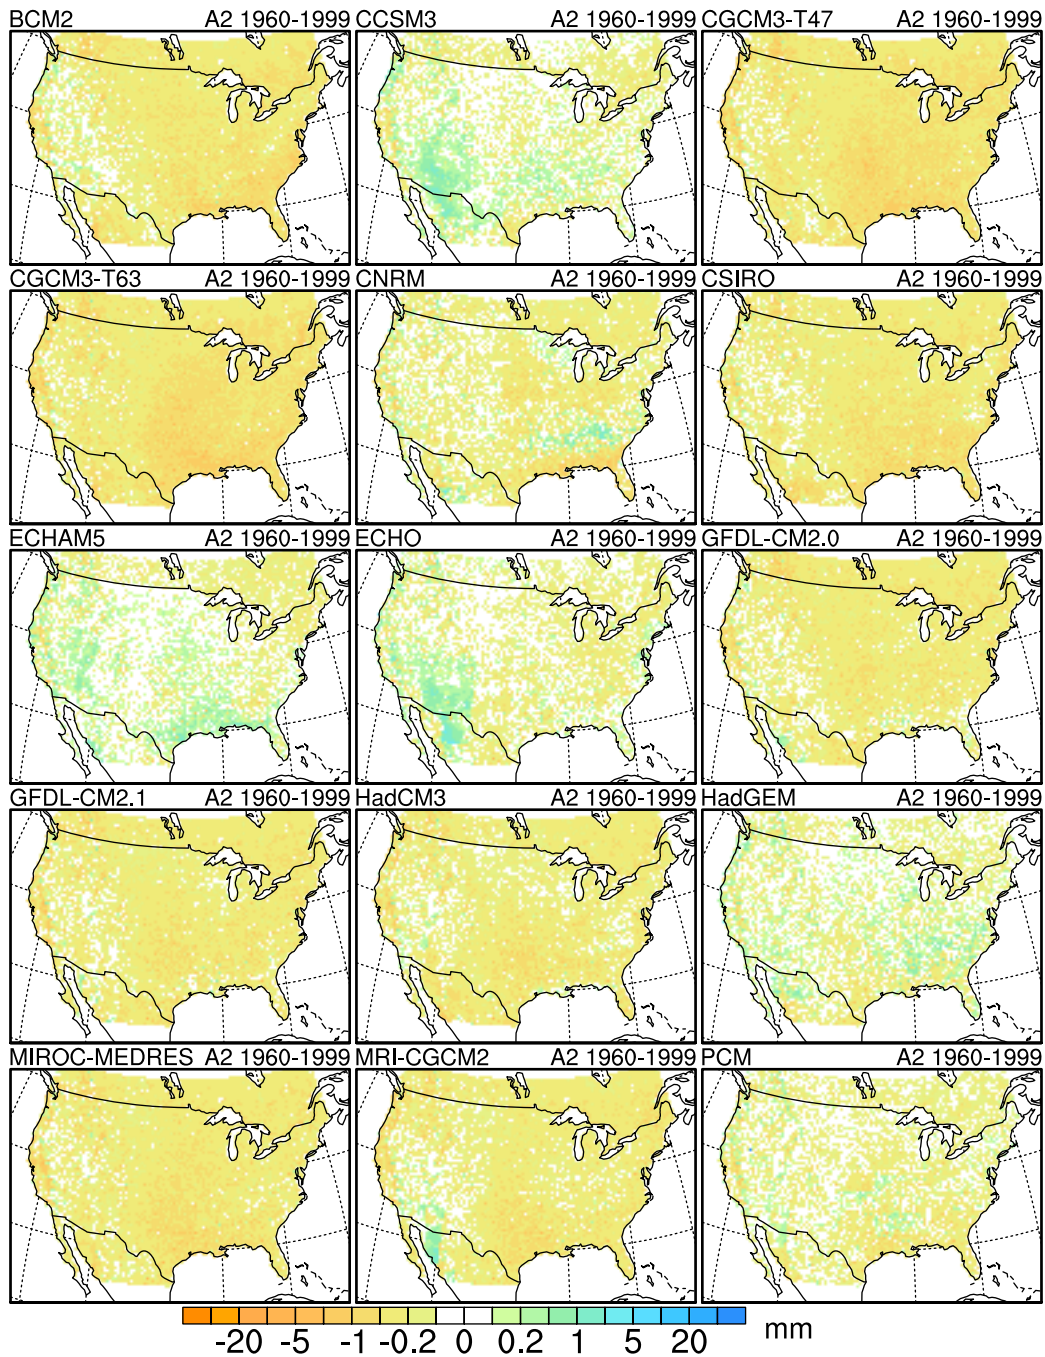
Pr Bias in 10th Quantile



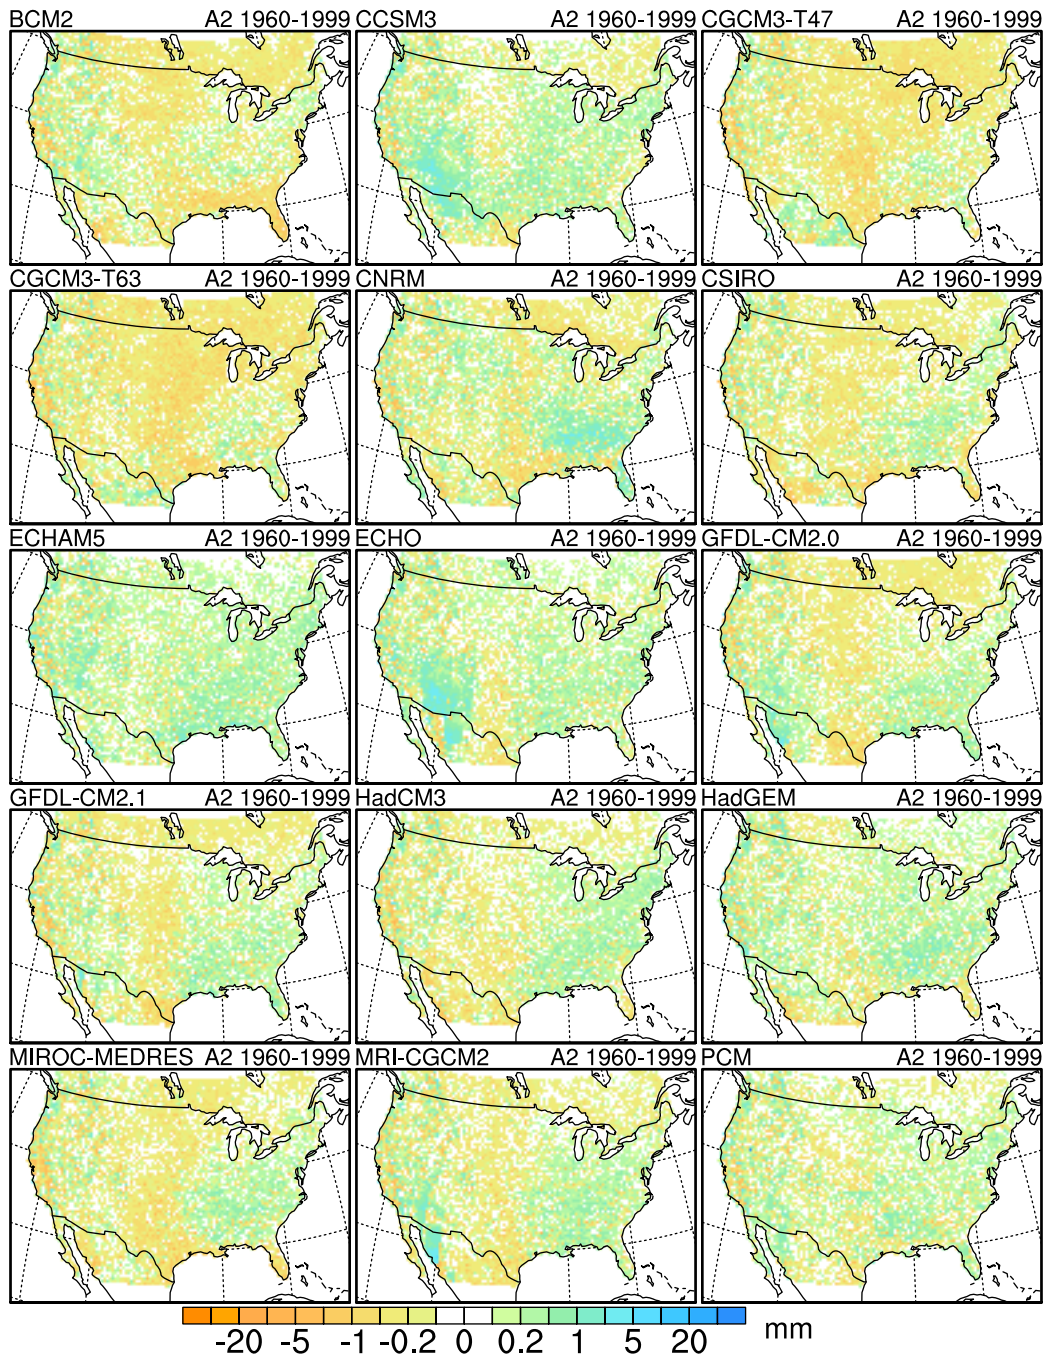
Pr Bias in 25th Quantile



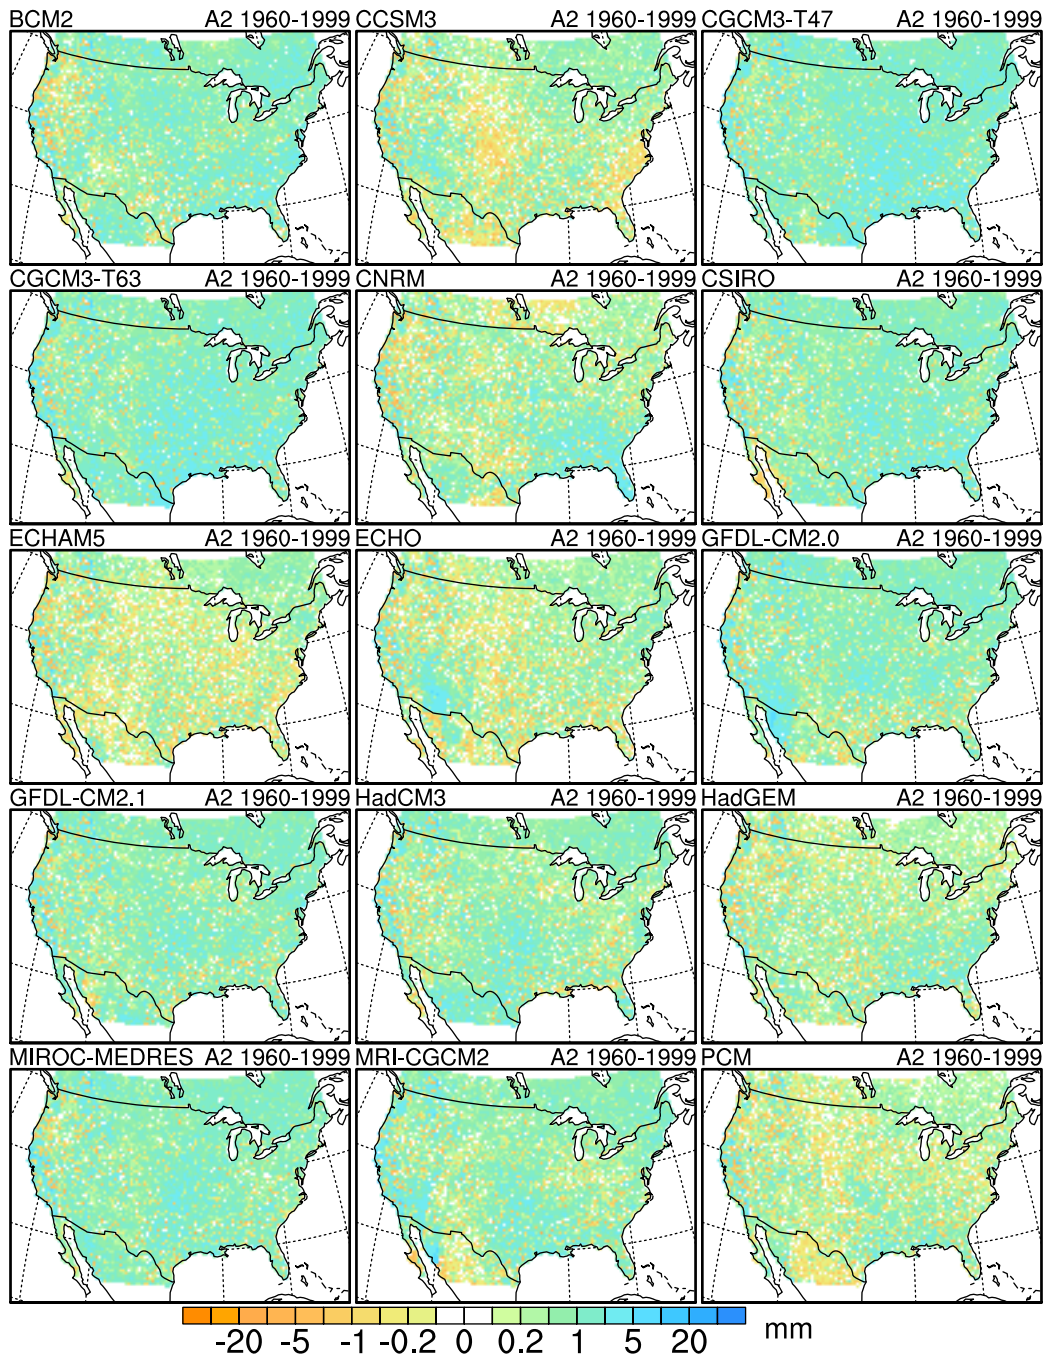
Pr Bias in 50th Quantile



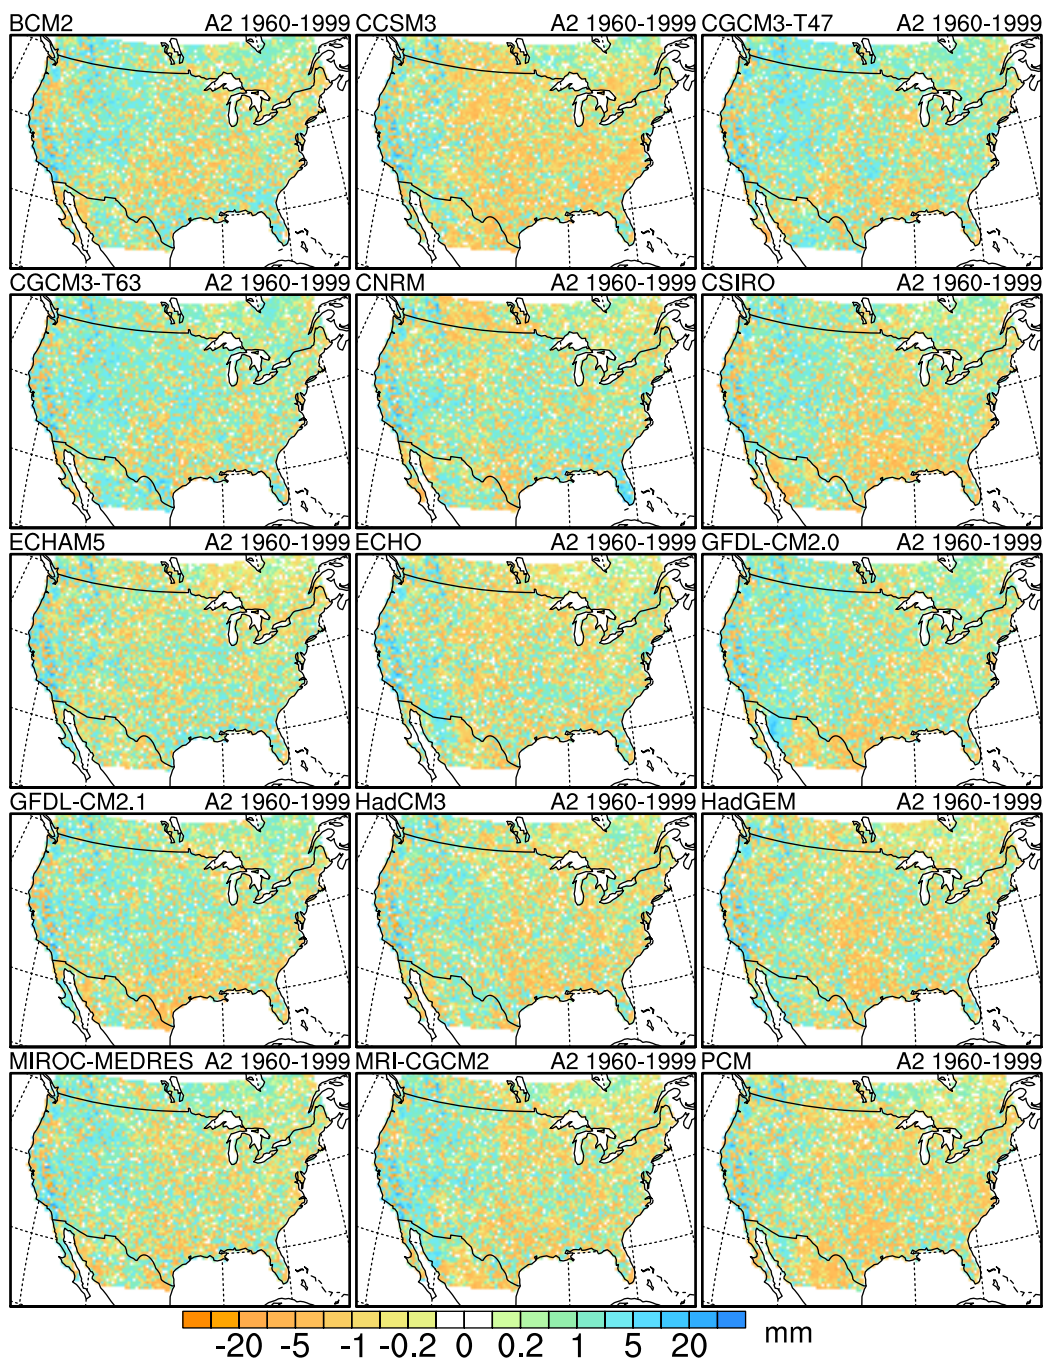
Pr Bias in 75th Quantile



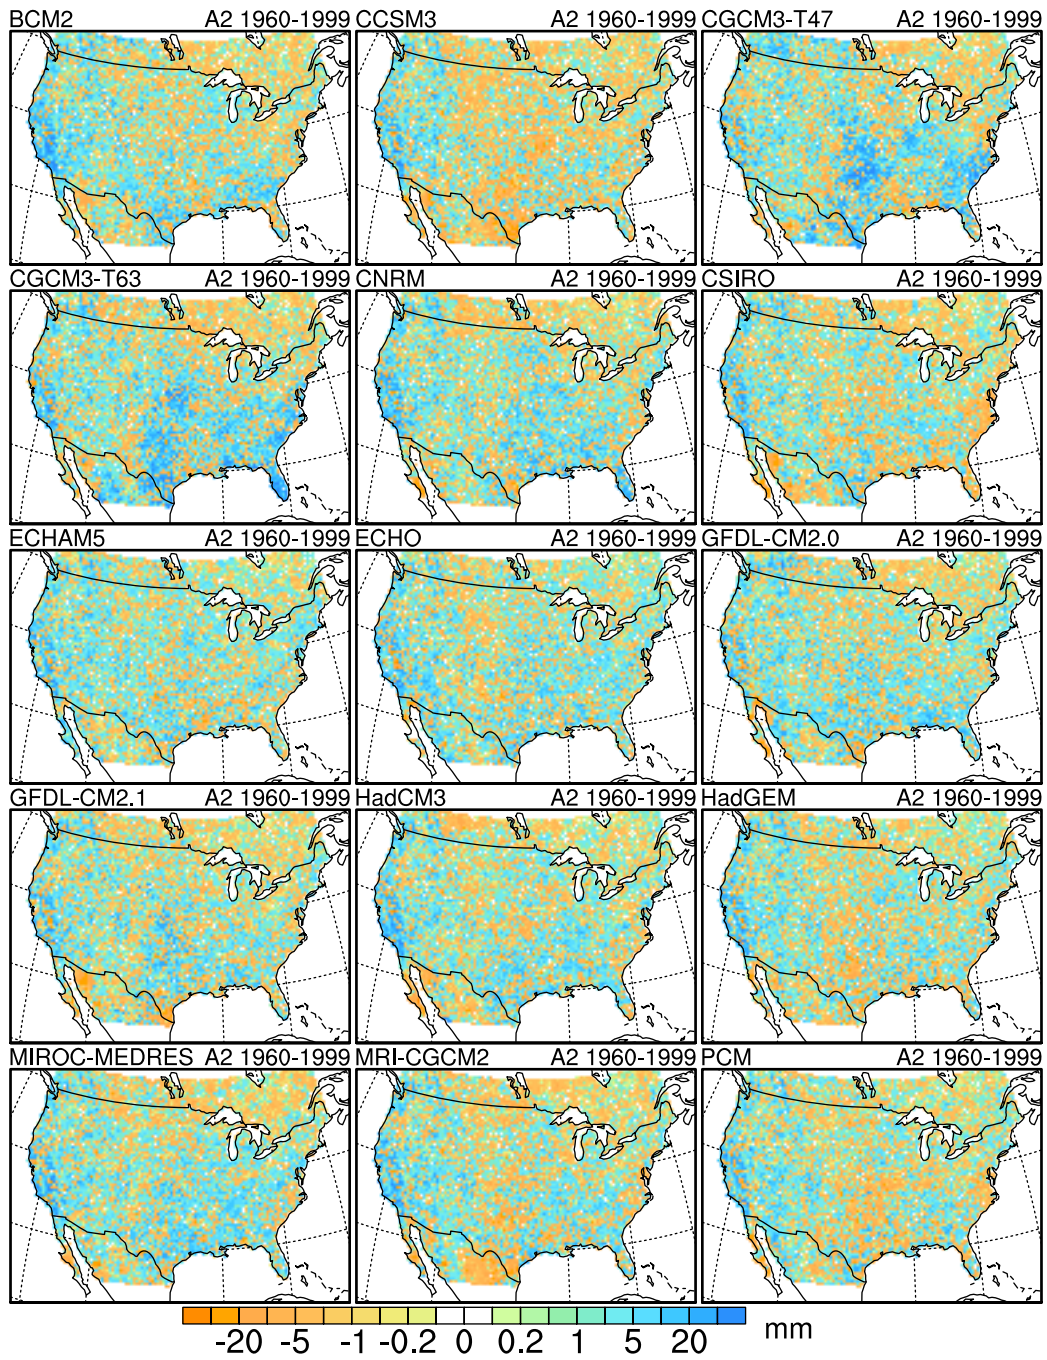
Pr Bias in 90th Quantile



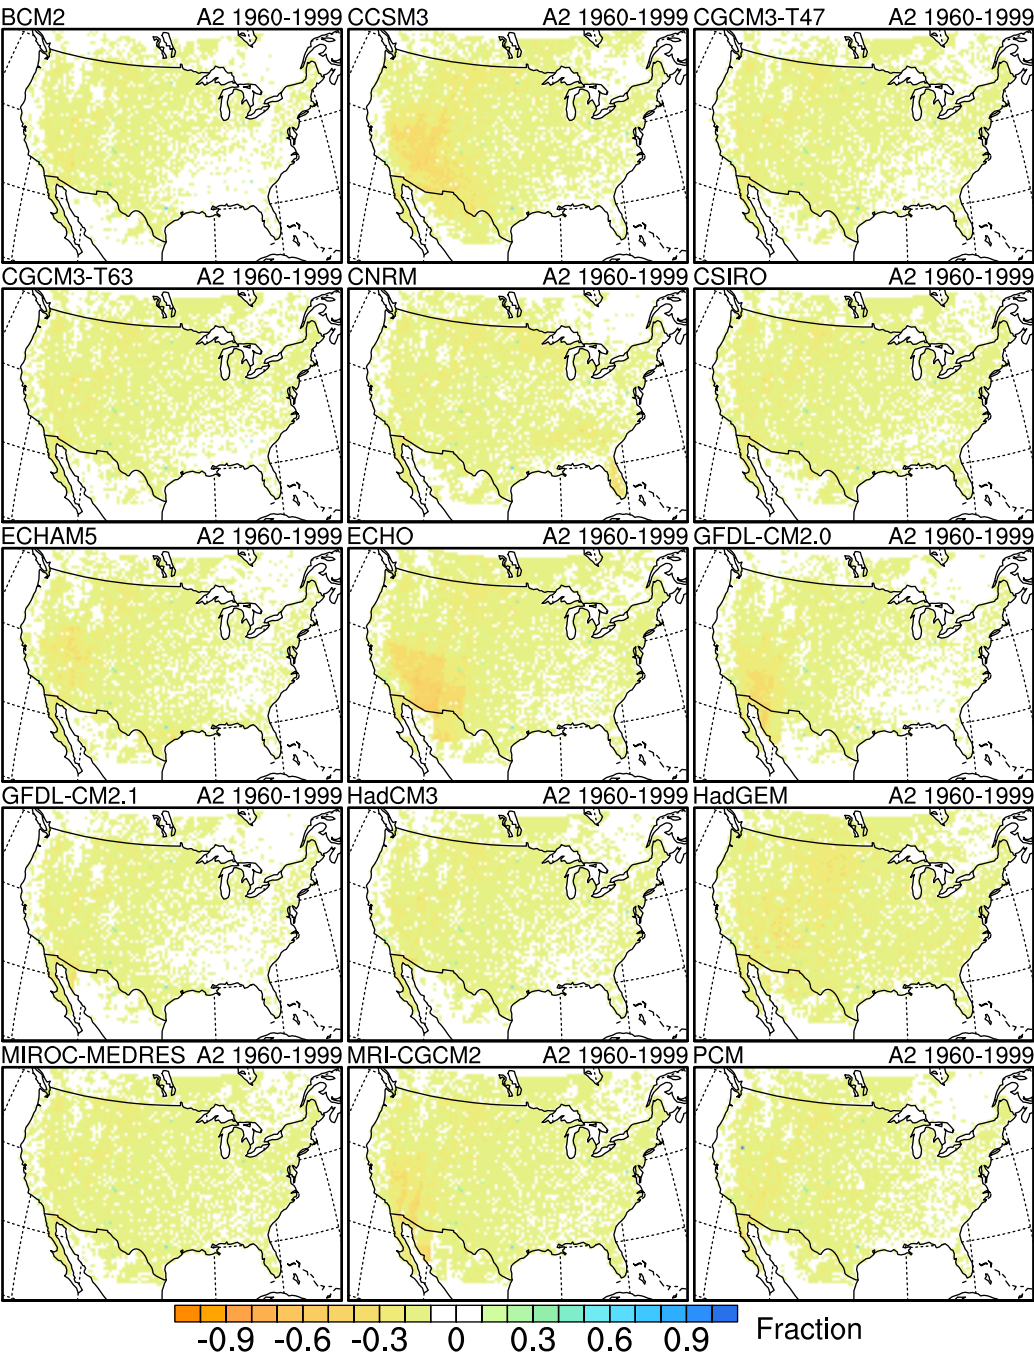
Pr Bias in 99th Quantile




Pr Bias in 99.9th Quantile




Bias in wet days with 24h cumulative Pr > 0 in



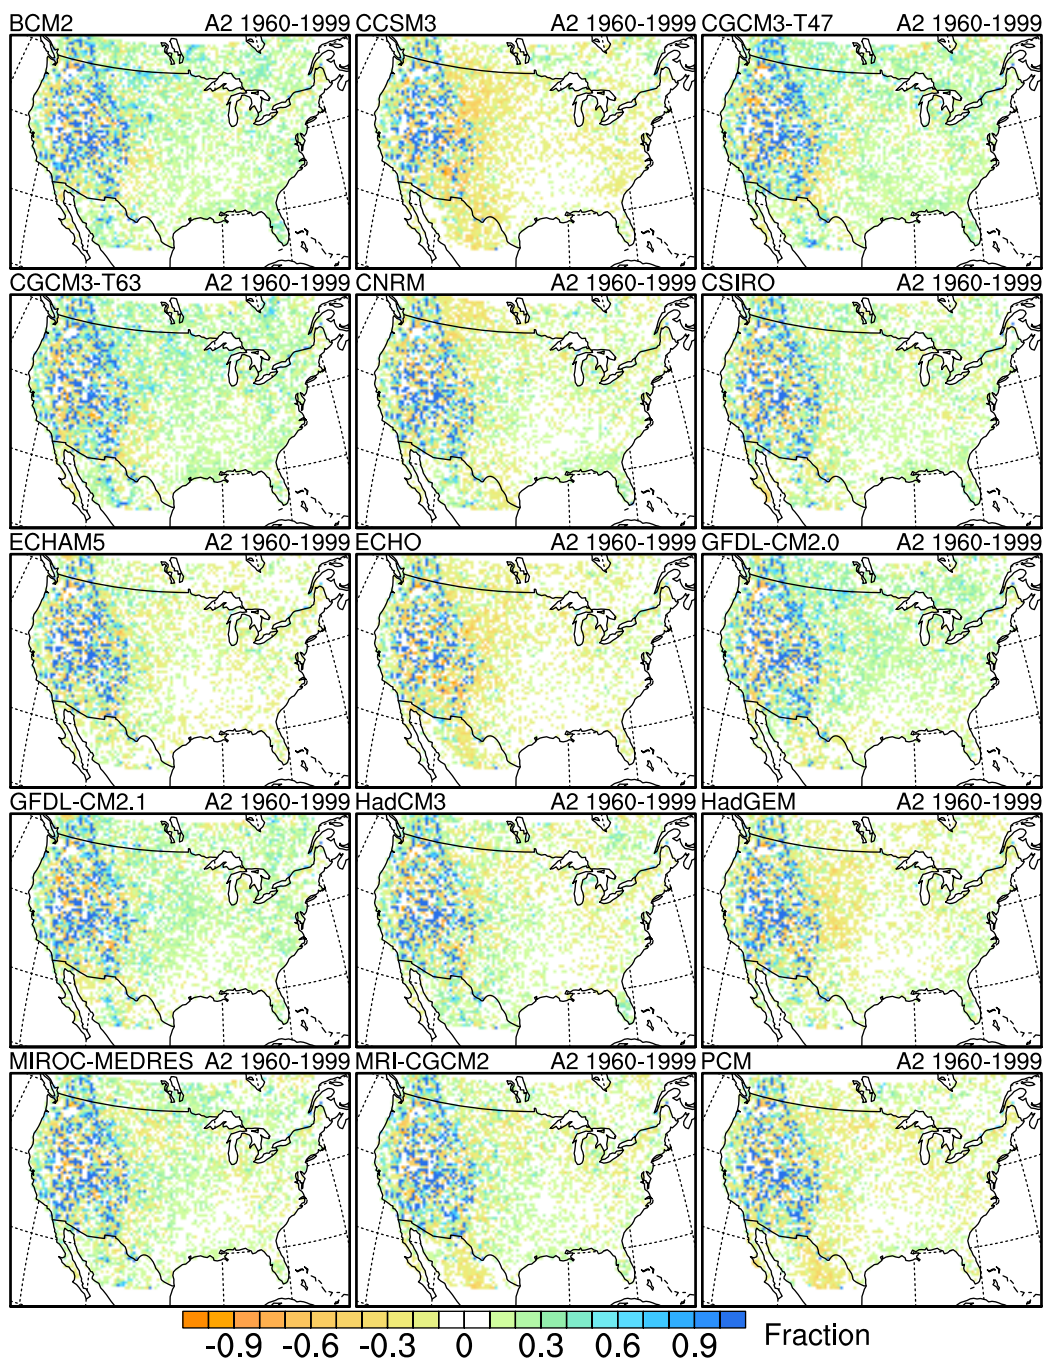
Bias in wet days with cumulative Pr > 0.01 in



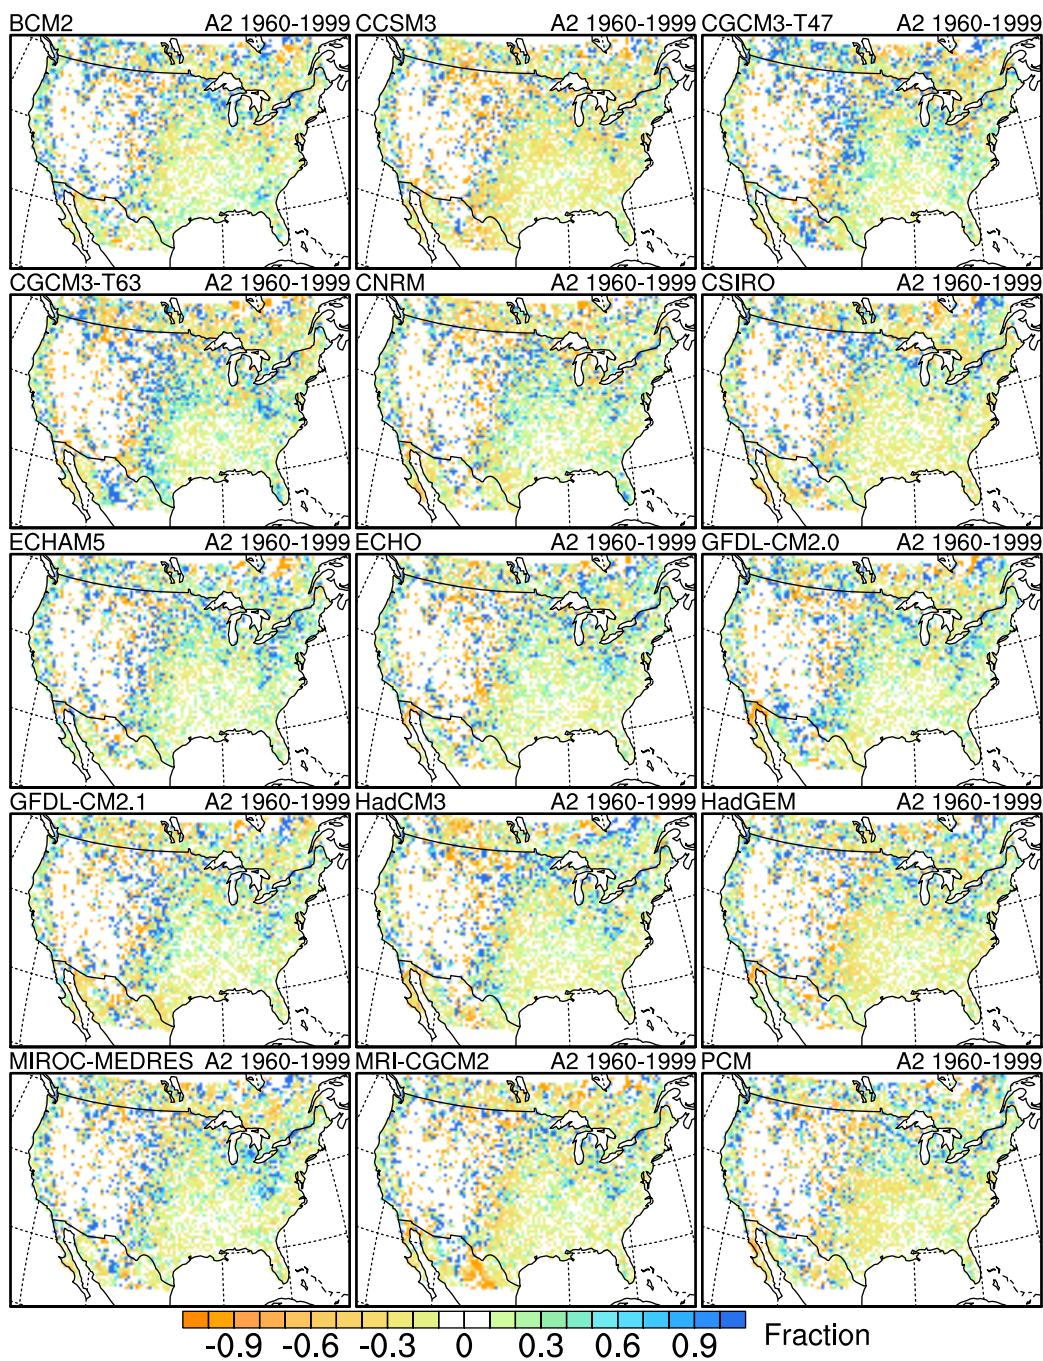
Bias in wet days with $Pr > 0.5$ in



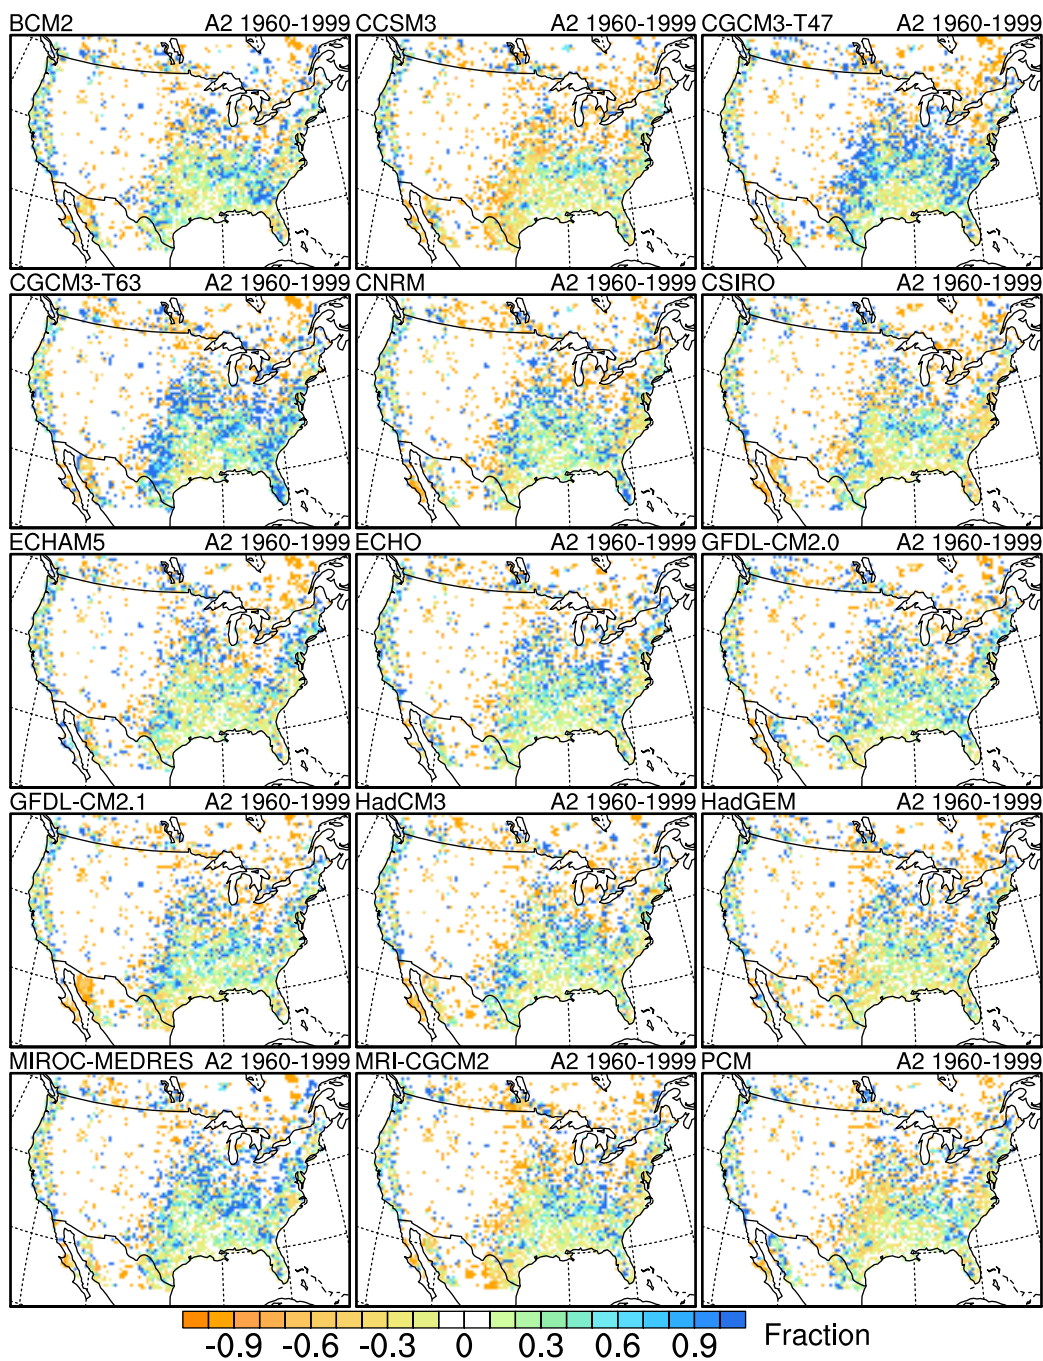
Bias in wet days with $Pr > 1$ in



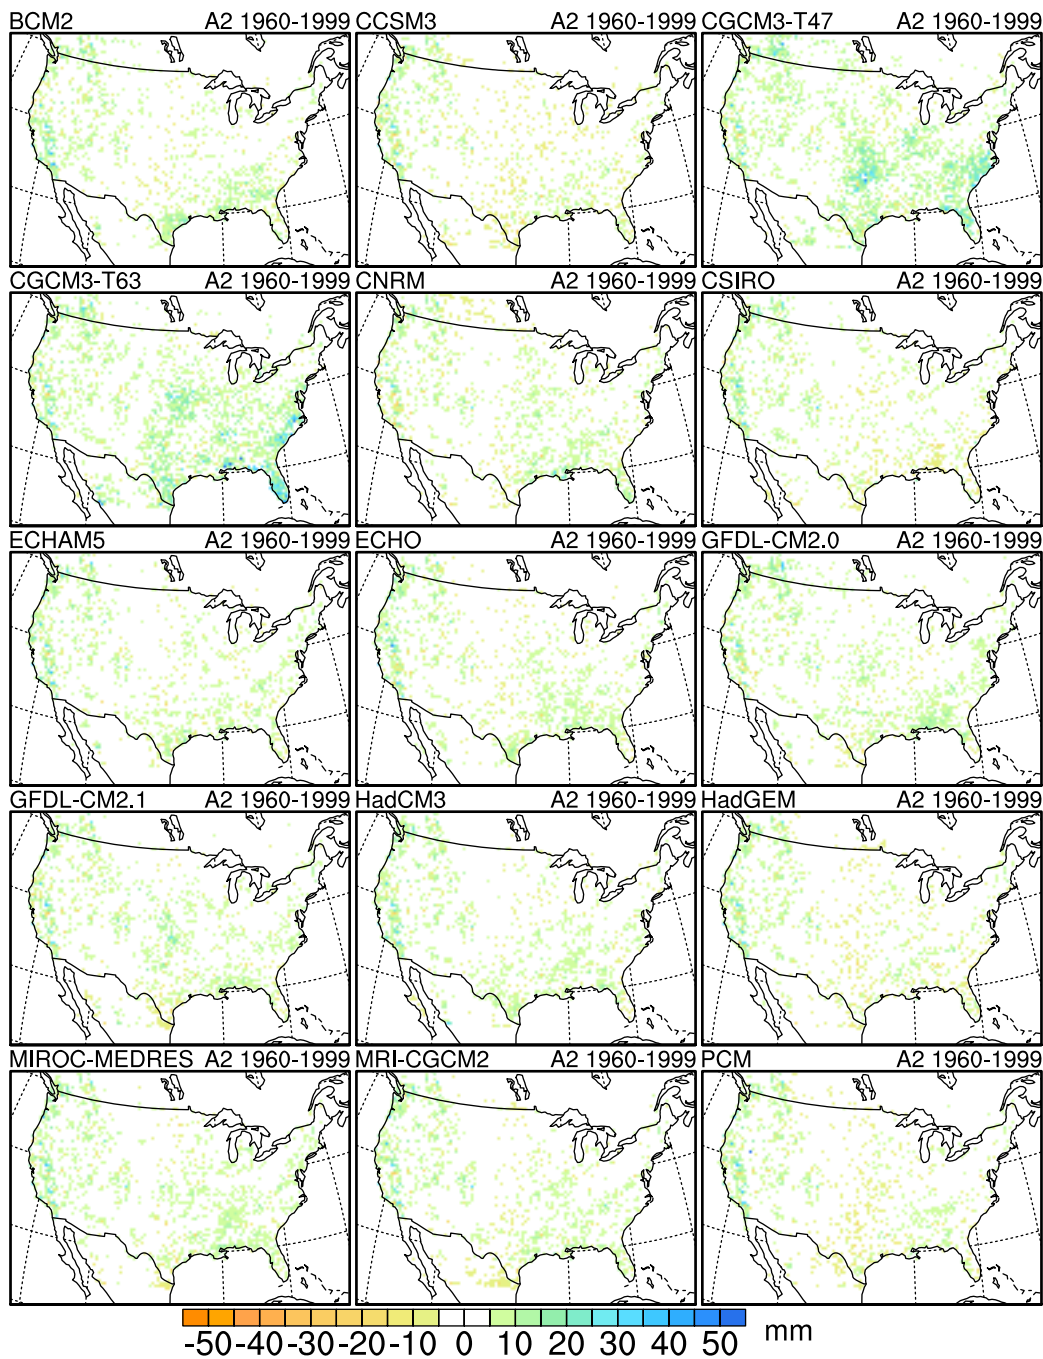
Bias in wet days with $Pr > 2$ in




Bias in wet days with $Pr > 3$ in




Bias in average 24h cumulative Pr on the wettest day of the year



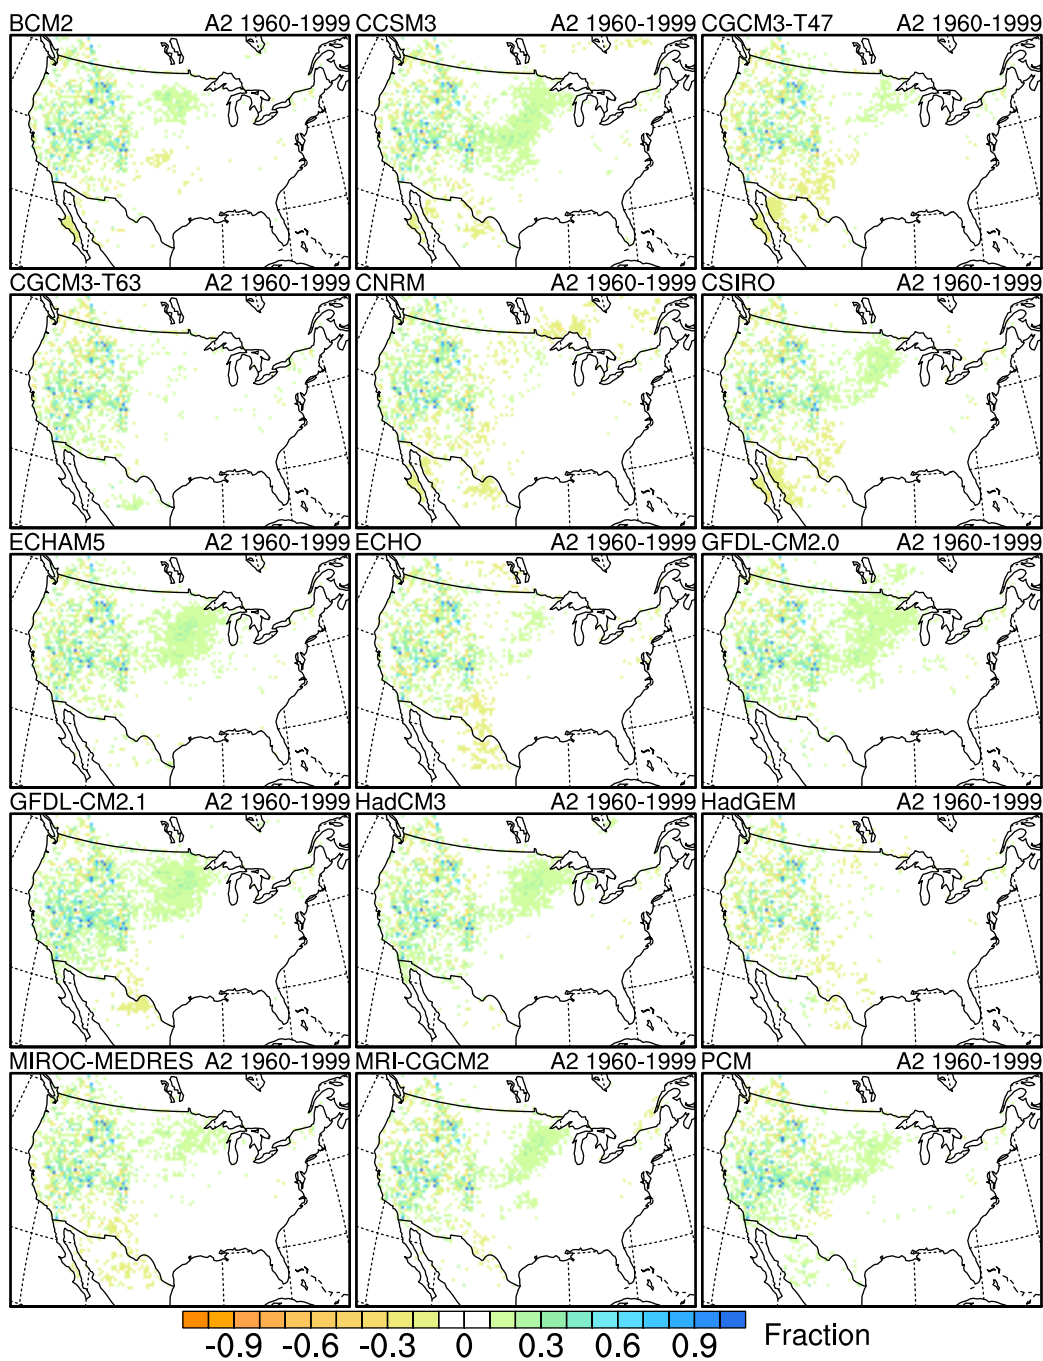
Bias in average 24h cumulative Pr on the 5-day wettest period of the year




Bias in 24h cumulative Pr on the wettest day in 30 years




Bias in winter cumulative Pr (DJF)




Bias in spring cumulative Pr (MAM)



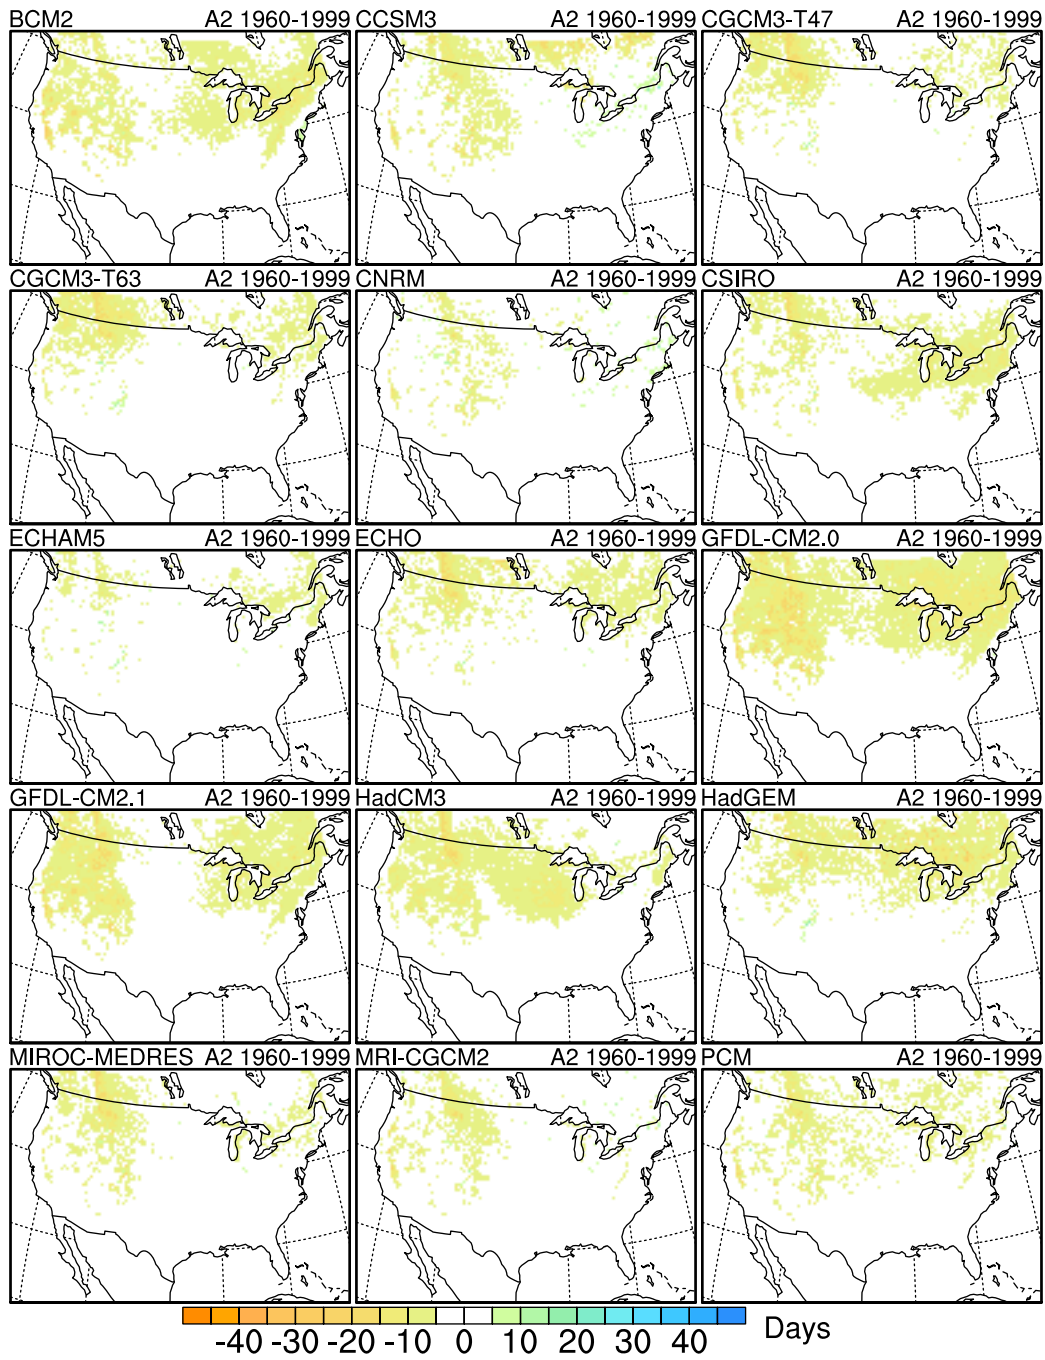
Bias in summer cumulative Pr (JJA)




Bias in fall cumulative Pr (SON)



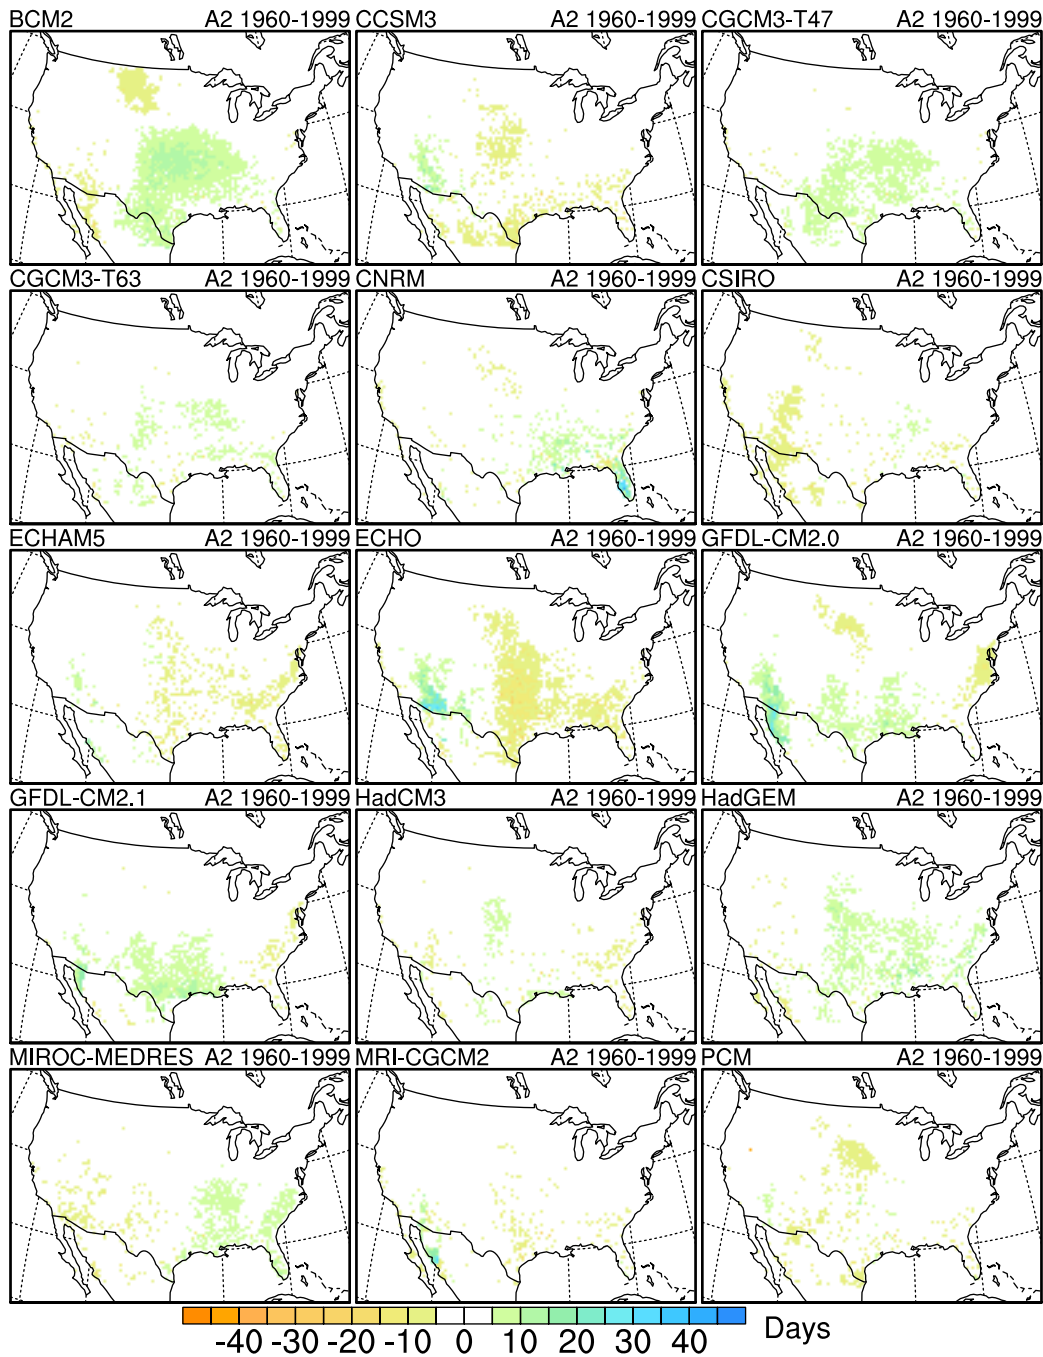
Bias in Number of snow days per year



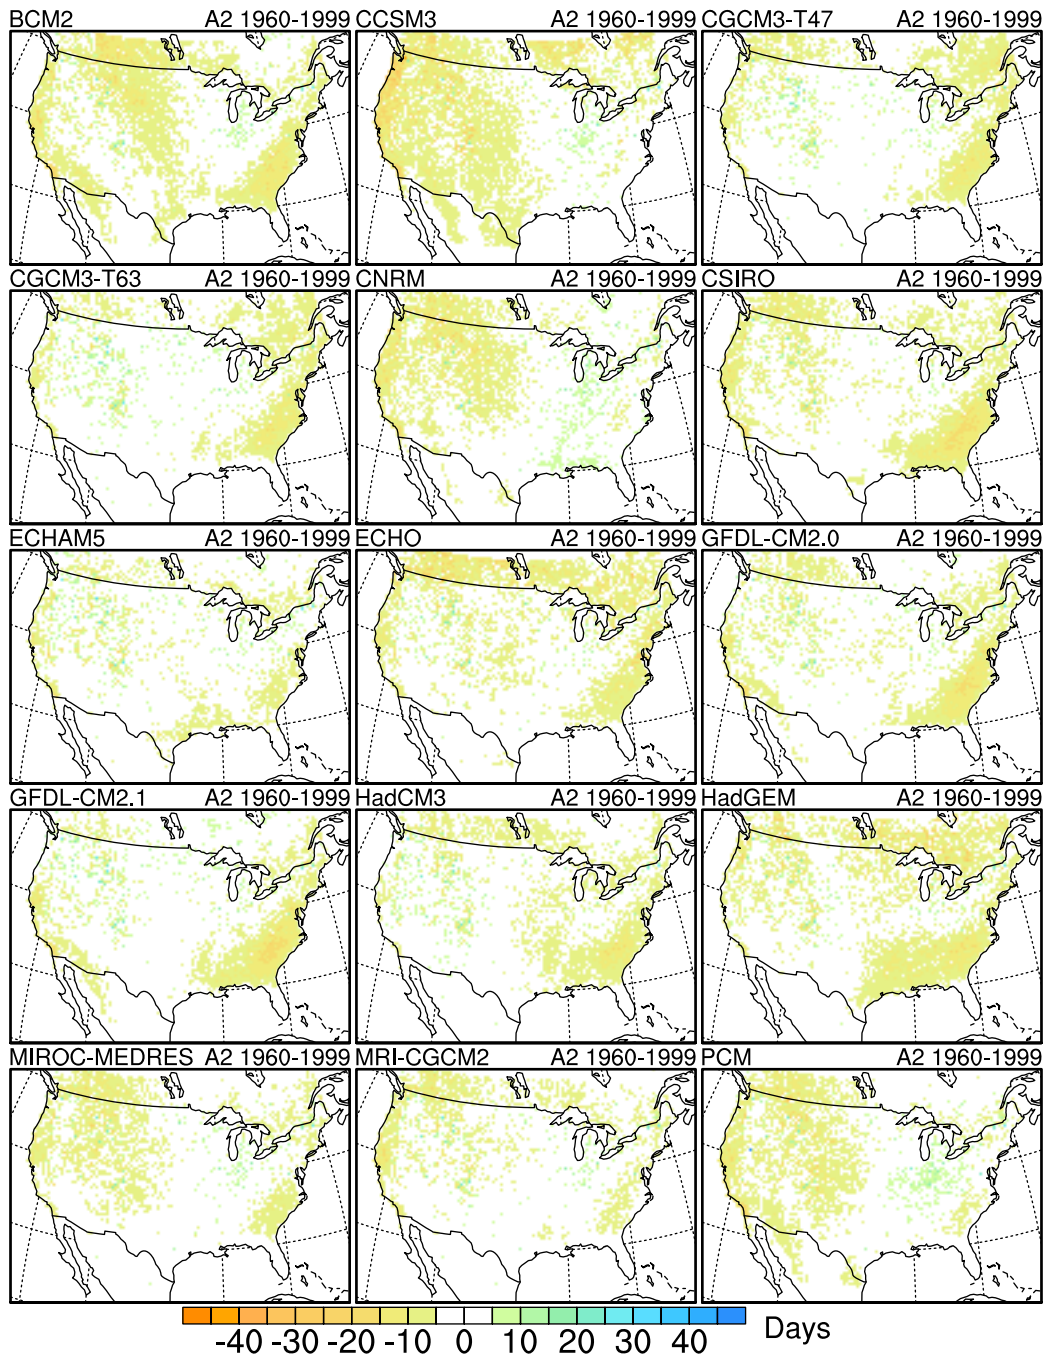
Bias in ratio of Pr falling as rain to that of snow



Bias in Number of hot, dry days per year ($Pr < 0.01in$, $T_{max} > 90F$)



Bias in number of cool, wet days per year ($pr > 0.01in$, $t_{max} < 65F$)



An asynchronous regional regression model for statistical downscaling of daily climate variables

Anne M. K. Stoner,^{a*} Katharine Hayhoe,^{a,b,c} Xiaohui Yang^b and Donald J. Wuebbles^d

^a Climate Science Center, Texas Tech University, Lubbock, TX, USA

^b ATMOS Research & Consulting, Lubbock, TX, USA

^c Department of Political Science, Texas Tech University, Lubbock, TX, USA

^d Department of Atmospheric Science, University of Illinois at Urbana-Champaign, Urbana, IL, USA

ABSTRACT: The asynchronous regional regression model (ARRM) is a flexible and computationally efficient statistical model that can downscale station-based or gridded daily values of any variable that can be transformed into an approximately symmetric distribution and for which a large-scale predictor exists. This technique was developed to bridge the gap between large-scale outputs from atmosphere–ocean general circulation models (AOGCMs) and the fine-scale output required for local and regional climate impact assessments. ARRM uses piecewise regression to quantify the relationship between observed and modelled quantiles and then downscale future projections. Here, we evaluate the performance of three successive versions of the model in downscaling daily minimum and maximum temperature and precipitation for 20 stations in North America from diverse climate zones. Using cross-validation to maximize the independent comparison period, historical downscaled simulations are evaluated relative to observations in terms of three different quantities: the probability distributions, giving a visual image of the skill of each model; root-mean-square errors; and bias in nine quantiles that represent both means and extremes. Successive versions of the model show improved accuracy in simulating extremes, where AOGCMs are often most biased and which are frequently the focus of impact studies. Overall, the quantile regression-based technique is shown to be efficient, robust, and highly generalizable across multiple variables, regions, and climate model inputs. Copyright © 2012 Royal Meteorological Society

KEY WORDS statistical downscaling; quantile regression; climate; temperature; precipitation

Received 29 April 2012; Revised 30 August 2012; Accepted 2 September 2012

1. Introduction

Atmosphere–ocean general circulation models (AOGCMs) and the new generation of earth system models provide insights into the dynamic nature of possible climate responses to anthropogenic forcing. With spatial scales typically on the order of one half degree or coarser, however, they are unable to simulate climate at the local to regional scale. To compensate for this relatively coarse resolution, a number of dynamical and statistical techniques have been developed to downscale climate model outputs to the impact-relevant spatial and temporal scales at which observations are made.

Despite the plethora of downscaling methods in the literature (Crane and Hewitson, 1998; Wilby *et al.*, 1998; Huth *et al.*, 2001; Stehlik and Bardossy, 2002; Wood *et al.*, 2004; Haylock *et al.*, 2006; Schmidli *et al.*, 2006; Kostopoulou *et al.*, 2007; Hidalgo *et al.*, 2008; to name just a few out of hundreds), relatively few downscaling methods have been applied to quantify potential impacts

of climate change at the local to regional scale for a broad cross-section of regions and sectors across North America. The majority of studies of climate change impacts in the United States, for example, rely on one of five methods: a delta approach whereby a change or ‘delta’ is added to observed mean annual, seasonal, or monthly values in order to get future values (Hay *et al.*, 2000; as used in USGCRP, 2000); simulations from a regional climate model (e.g. Mearns *et al.*, 2009; as used in NARCCAP); the Bias Correction-Statistical Downscaling model originally developed as a front end to the hydrological variable infiltration capacity model, which uses a quantile mapping approach to downscale monthly AOGCM-based temperature and precipitation to a regular grid (Wood *et al.*, 2004; as used in Hayhoe *et al.*, 2004, 2008; Luers *et al.*, 2006; USGCRP, 2009); a constructed analogue approach that matches AOGCM-simulated patterns to historical weather patterns (Hidalgo *et al.*, 2008; as used in Luers *et al.*, 2006); and a linear asynchronous regression approach that downscales daily AOGCM-based temperature and precipitation to individual station locations (Dettinger *et al.*, 2004; as used in Hayhoe *et al.*, 2004, 2008, 2010).

Each of these methods has its own benefits, and each can be sufficient for certain applications. For example,

* Correspondence to: A. M. K. Stoner, Climate Science Center, Texas Tech University, 113 Holden Hall, Boston & Akron Streets, Lubbock, TX 79409-1015, USA. E-mail: anne.stoner@ttu.edu

the simple and transparent delta approach can yield a nearly identical downscaled annual or seasonal mean temperature value as a more complex statistical model. At the other end of the spectrum, complex regional climate models are computationally demanding, but provide consistent high-resolution projections for a plethora of surface and upper-air variables. None of these five methods, however, allows for using multiple climate models and scenarios as input while downscaling to any spatial scale (including both station-based and gridded), simulating additional impact-relevant variables (such as solar radiation and humidity), and adequately resolving projected changes in daily climate extremes, at the same time.

For that reason, we have developed a new statistical downscaling model, the asynchronous regional regression model (ARRM). ARRM builds on the same statistical technique used by the last downscaling approach listed above (Dettinger *et al.*, 2004), asynchronous quantile regression, to define a quantitative relationship between any daily observed and simulated surface variable that has a symmetric distribution, with particular emphasis on accurately resolving the relationship at the tails of the distribution in order to capture simulated changes in extremes. Asynchronous quantile regression removes the time stamp from historical observations and simulations, reordering each time series by value before matching quantiles of observed data with those from AOGCM output. This is important because coupled AOGCM simulations generate their own patterns of natural variability, meaning that no day-to-day or even year-to-year correspondence with observations should be expected.

The general concept of quantile regression was originally developed in the field of econometrics by Koenker and Bassett (1978) to estimate conditional quantiles of the response variable as opposed to the conditional mean estimated by the orthodox least-squares regression method. The quantile regression approach is of particular utility to geospatial data, in that it can be used to determine relationships between two quantities that are not measured simultaneously, such as an observed and a model-simulated time series. It takes advantage of the hypothesis that although the two time series may be independent, their distributions may be similar.

The general technique of quantile regression has been used in a variety of applications, including by O'Brien *et al.* (2001) to determine relationships between measurements of relativistic electron conditions measured from two different satellites passing over the same area at different times. Dettinger *et al.* (2004) were the first to apply this method to downscaling AOGCM output, to examine simulated hydrologic responses to climate change. In this application, the first time series was observations and the second, historical model simulations. The regression model derived from these two distributions was then applied to transform the distribution of, or downscale, future model simulations.

The objective of this study is to build on the foundation of quantile regression to develop a relatively straightforward, flexible, efficient, and robust statistical model

that is capable of downscaling any atmospheric variable, measured on a daily or monthly basis, which has, or can be transformed into, an approximately symmetric distribution. Section 2 describes the statistical basis of the model and refinements that improve its ability to downscale global model outputs. Section 3 describes the long-term weather station observations and the AOGCM outputs used to evaluate the downscaling model in terms of its ability to simulate observed temperature and precipitation, using the same variables from the AOGCMs as predictors. Section 4 describes how the model was developed in multiple steps, each of which is successively tested to ensure that the additions improve the model's ability to reproduce historical climate. Section 5 discusses the results of applying the downscaling model to end-of-the-century temperatures and precipitation and the changes between downscaled and raw AOGCM output compared with present conditions. Finally, Section 6 summarizes the findings of this study.

2. Model development

2.1. Statistical basis

The concept of *quantile regression* was first introduced by Koenker and Bassett (1978), where quantiles refer to values of a cumulative population (i.e. when the data are sorted by increasing value) that divide the population into equal-sized segments. Quantiles are the data values marking the boundaries between consecutive subsets. If the data are divided into q equal-sized subsets, the k th quantile for a variable is the value x such that the probability that the variable will be less than x is no greater than k/q and the probability that the variable will be more than x is no greater than $(q - k)/q$. A distribution has $q - 1$ quantiles, one for each integer k satisfying $0 < k < q$.

In general, regression analysis quantifies covariance between variables, and, if it exists, provides a model to predict one variable on the basis of the other variables used as input to the regression. Quantile regression specifically estimates *conditional quantile functions* – models in which quantiles of the distribution of the predictor variable are expressed as functions of observed covariates (Koenker and Hallock, 2001). In other words, quantile regression results in estimates approximating the quantiles of the predictor variable. For a time series containing N values there are N ranks in each vector. A model can be constructed by regressing the value at rank n_i of the simulated vector onto the value of the same rank of the vector containing observed values, for $i = 1 \dots N$ (as done for example in Dettinger *et al.*, 2004). This regression is asynchronous, i.e. data values that are regressed against each other did not necessarily occur the same calendar day, but rather correspond by quantile or rank. The regression model derived from historical AOGCM simulations and historical observations can then be applied to future AOGCM simulations, to project downscaled future conditions.

Asynchronous regression is an important component of this model, because a coupled AOGCM simulation is free to evolve chaotically, with only the external forcings being prescribed; hence, each simulation represents one out of many possible outcomes and no daily correspondence between the model and observations should be expected.

2.2. Model input

Both theoretical and practical considerations affect the selection of inputs to quantile regression. First, it is important to verify that the two time series (simulated and observed, or predictor and predictand) have somewhat similar distributions; the closer both distributions are to Gaussian, the simpler the function relating the two distributions. Even non-Gaussian distributions can sometimes be manipulated to mimic a Gaussian distribution; here, in the case of precipitation, by taking the natural logarithm of the daily wet day precipitation values.

To train the downscaling model, the observed record must have an adequate length and quality of data. A minimum of 20 consecutive years of daily observations with less than 5% missing data is usually needed in order to appropriately sample from the range of natural climate variability at most of the station locations examined here and to produce robust results without overfitting. To challenge the downscaling model, two stations were selected for this evaluation that had substantially less data available (Bridgeport, WV with 78% and Moosehead Lake, ME with 88% of daily data missing over 50 years).

2.3. Model structure

The structure of the ARRM model is summarized in Figure 1. The first step is to prepare the data by separating it into 12 vectors by month such that a separate statistical model can be built for each month. This accounts for different weather patterns dominating any given region at different times of the year that could alter AOGCM biases relative to observations. Two weeks of overlapping data on either side of each month are included to account for future conditions that may lie outside the range of a typical historical month. This extension also doubles the use of each data point during the training process. Each month's time series is then reordered by rank to create an asynchronous vector. Figure 2 shows AOGCM-simulated (grid cell containing the weather station) versus observed temperature for chronological and for sorted data, illustrating how ranking of the inputs provides a correlation between observations and model simulations whereas matching by calendar date does not.

The second step in the ARRM model is to fit a regression function to the ranked values shown in Figure 2(b). For most station locations and global models, a linear fit (as used in Dettinger *et al.*, 2004) is adequate within at least the 20th–80th percentiles of the distribution (dark-coloured line in Figure 3) with a high coefficient of determination (R^2). However, residuals are often large

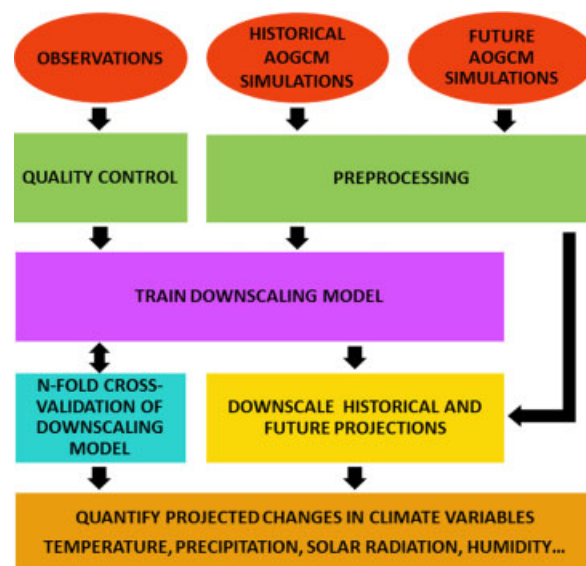


Figure 1. Structure of the ARRM downscaling model. This figure is available in colour online at wileyonlinelibrary.com/journal/joc

near the tails of the distribution that, depending on the application, can be of greater interest to climate impact studies than values at the centre of the distribution. Polynomials of increasing order result in increasingly better fits to the historical observations (not shown), but run two serious risks: first, of overfitting, and second, of exhibiting unnatural behaviour at the tails of the distribution that could unrealistically predict lower observed temperatures for higher modelled values than for lower modelled values, and vice versa.

Instead, we found that a piecewise linear regression (light-coloured, segmented line, Figure 3) provided the most consistent fit while accounting for biases in model values near the tails of the distribution; biases that can be markedly different than those simulated for values near the centre of the distribution. Adding breakpoints, or knots, allows for different slopes at different parts of the distribution, in particular minimizing the residuals at the tails of the distribution when compared with either a linear or a polynomial fit.

R (R Development Core Team, 2012), the statistical programming language used to build ARRM, has spline-based functions such as *bs* and *ns* that can add breakpoints to a regression. However, these functions require the user to set the number of breakpoints manually and then place the points at predetermined, evenly distributed, quantiles. As illustrated in Figure 3, the ideal number of breakpoints can vary broadly, depending on the characteristics of model bias for a given month and/or location. A new function was therefore required that would optimize the regression model for each month by automatically identifying the number and location of up to six independent breakpoints. This piecewise linear regression function is described next.

The third step in the ARRM model is to use the statistical regression models, constructed from observed and historical simulated time series, to downscale future




Figure 2. Monthly scatter plot of observed and AOGCM (CCSM3) simulated (nearest grid cell) daily maximum 1960–2009 January temperature for Bridgeport, West Virginia, matched (a) by calendar date and (b) by rank.




Figure 3. Scatter plot of observed *versus* AOGCM (CCSM3) simulated daily maximum 1960–2009 January temperature for Bridgeport, West Virginia (including half of previous and following months), ranked by quantile. Dark-coloured line shows the results of a linear regression on the data; lighter, segmented line shows the results of a piecewise linear regression; and vertical grey lines identify knots in the piecewise linear regression. This figure is available in colour online at wileyonlinelibrary.com/journal/joc

projections. The resulting downscaled values must subsequently be rearranged back into the original order to retrieve the final product, a continuous chronological time series of the downscaled values.

2.4. Piecewise linear regression function

The piecewise linear regression function developed for ARRM is based on linear regression that iterates over a moving window. For the majority of the distribution, the window width remains fixed at a given percentage of the total number of data points for that particular month. As the concentration of data points near the tails of the distribution is much sparser than at the centre, window

width at the tails of the distribution decreases linearly to a minimum width by the ends of the distribution.

This function requires four fixed settings: the percentage of data points in the fixed window width, the minimum and maximum probabilities over which a fixed window width is used, the minimum permissible width of the window at the tails of the distribution, and the maximum number of breakpoints allowed. Optimal values for these settings are a function of AOGCM bias, characteristics of which differ from one variable to another. In general, a fixed window width of 5% of the distribution between probabilities of $0.1 < P < 0.9$, linearly decreasing to a minimum width of either 2 °C or ten data points (whichever is greater) for $P < 0.1$ and $P > 0.9$, is adequate for temperature as the relationship between observed and modelled values tends to be relatively linear over much of the distribution. For precipitation, greater variability in AOGCM bias over the distribution requires a wider fixed window width, on the order of 10%, between probabilities of $0.15 < P < 0.85$, linearly decreasing to a minimum of 5% of the mean value or ten data points (again, whichever is greater) for $P < 0.15$ and $P > 0.85$.

Up to six breakpoints are allowed in each regression model. This number was determined on the basis of two factors: first, visual testing by plotting downscaled projections for the historical period for individual months showed that more breakpoints tended to increase the risk of overfitting, such as introducing shorter segments with negative slopes, particularly for months with sparse data or poor model performance, and second, months with dense data rarely required more than six breakpoints and often far less. The function begins the piecewise regression at the 40th percentile, where the data point at the 40th percentile is the largest value in the window and moves up (to the highest quantile of the distribution) from there. In other words, the moving window starts with the $X\%$ data points below the 40th percentile, where X equals 5 for temperature and 10 for precipitation. The selection of the 40th percentile is to ensure that

the middle part of the distribution is well covered by a moving window. QR matrix decomposition is used to fit a linear regression to the data in the window. The R^2 value for each regression is recorded and saved in a vector, and the moving window is shifted up one data point towards the end of the distribution until it reaches the 100th percentile. The first breakpoint is defined as the central point of the window with the lowest R^2 value of the vector, if the value of R^2 for that window is less than the value for the entire time series. The R^2 values on either side of that breakpoint are then blocked for the width of the moving window and a new minimum identified, for a total of up to three breakpoints in the upper half of the distribution.

This process is repeated beginning at the lowest found breakpoint, or if no breakpoints are found, at the 40th percentile moving down to the 0th percentile. This time, the moving window trails above the percentile. This allows an R^2 value to be assigned to each data point in the monthly vector, from the first to the last. Setting a minimum window width of ten points means that breakpoints are not allowed to fall within the first and last five points of the dataset.

Before the statistical model is finalized, slopes between breakpoints are automatically reviewed. Breakpoints that create a negative slope can cause lower AOGCM values to produce higher downscaled values than higher AOGCM values. Breakpoints that create a slope close to zero ($-0.1 < \text{slope} < 0.1$) can create an unrealistic peak of nearly identical values in the downscaled distribution. Removal of a breakpoint causing a negative or 'flat' slope will always have a detrimental effect on the R^2 value of the regression fit, because the segment having the negative or 'flat' slope yielded the best fit, but improve the realism and generalizability of the fit. Sometimes, when AOGCM biases are particularly nonlinear, the removal of negative slopes can have a greater impact on the quality of the fit than the impact of having a few data points with downscaled values that decrease rather than increase for a small interval within the distribution. Hence, the function allows for negative or 'flat' slopes under two conditions: if they are not the first or last segment in the regression, and if they span less than five points. If these conditions are not met, the breakpoint below the negative slope is removed unless it is the first segment of the regression, in which case the breakpoint to the right is removed. One breakpoint is removed at a time and the process repeated once the regression and new slopes have been recalculated to determine whether a new segment with a negative slope has been introduced. This process is repeated until all negative or flat slopes have been eliminated.

Once the breakpoints have been finalized, the regressions are used to build a statistical model that performs piecewise linear regressions, with the use of spline interpolations, between the monthly simulated and observed data ordered by rank. This regression model can then be used to downscale future values, similarly ordered by rank, assuming stationarity in climate system feedback mechanisms.

2.5. Bias correction

As ARRM is a statistical model, there is a risk of introducing unrealistic values especially at the tails of the distribution, where data points are sparse and the slope of the initial and/or final regression can be very sensitive to a single extreme point. In some cases, an observational data point may even be in error. An example is the Global Historical Climatology Network (GHCN) dataset for Hialeah, FL, which had a recorded maximum temperature for 8 November 2003 with a value of -17.8°C , 25°C lower than the second lowest maximum temperature recorded for this station, and with temperatures for the previous and following days of 29.4 and 30.0°C , respectively. This erroneous point noticeably affected the magnitude of predicted cold temperature extremes for this location. Unrealistic values in the original observations are therefore removed by the quality control procedure described in Section 3, prior to their use as input to the downscaling model.

Because of this sensitivity, downscaled extremes (defined as lying below the 5th percentile and/or above the 95th percentile of the distribution) that fall outside a realistic range for each station are corrected separately, by calculating the bias in percent difference between the downscaled value and the minimum or maximum observed value for that location. To avoid large biases that can be caused by small differences between low values, temperature is first converted to Kelvin and an arbitrary large number (here, 250) is added to daily precipitation values. For temperature, scaling is done by dividing the downscaled value by $1 + \text{the bias}$ when values fall more than 3% below or above the lowest or highest observed values (in Kelvin), respectively, or more than 2% above the highest observed precipitation value (with 250 added). For precipitation, the downscaling model in some cases predicts values below zero. These are reset to zero.

2.6. Variable-specific refinements

Although the downscaling model is purposely designed to be applicable to any variable with a relatively symmetric distribution, predictors must be preselected for each variable and there are some differences in the initial processing of each predictor that can improve the performance of the model in downscaling.

Selection of predictors for temperature and precipitation downscaling has been the subject of several comparative studies (Huth, 1999; Wilby and Wigley, 2000; Widmann *et al.*, 2003; Jeong *et al.*, 2012). ARRM has been designed to allow for user-selected predictors, if desired. For the purposes of model evaluation and comparison, predictors were chosen to be the same variables as the predictands: 2 m maximum and minimum temperature and 24 h cumulative precipitation. These are the most frequently archived daily output from both CMIP3 and CMIP5 AOGCMs; furthermore, comparison with upper-air predictors for the stations in this study showed no consistent improvement that would affect the performance of

the downscaling model. For models that archive convective, total, and/or large-scale precipitation, the downscaling model calculates the RMSE for the historical training period between the observations and separate downscaled values using each of the three predictors. The predictor variable and corresponding regression model for the training period with the lowest RMSE for a particular month is used to downscale future precipitation for that month and station. This refinement significantly improved the method's ability to simulate precipitation over regions that tend to experience more convective-type precipitation, including the subtropics and mid-latitude summer.

Smoothing AOGCM output has been previously recommended (e.g. Raisanen and Ylhäisi, 2011), and it has been suggested that the smoothing that results from averaging may be one of the reasons why ensemble AOGCM projections typically outperform any individual model simulation (Knutti *et al.*, 2010). Here, temperature fields are smoothed using Empirical Orthogonal Function (EOF) analysis, retaining only the EOFs accounting for 97% of the original variance. Root-mean-square errors (RMSEs) identified 97% as a generally appropriate threshold, with both higher and lower thresholds resulting in higher errors. This step improved model performance, especially for inland stations with higher variance.

Compared to temperature, precipitation tends to display a greater amount of smaller scale variability. This is likely one of the reasons why EOF filtering was found to degrade rather than assist precipitation downscaling. Precipitation is also a combination of a binary (wet/dry) and a continuous non-Gaussian distribution that must be transformed into a more symmetrical distribution before it can be ranked by quantile. Dettinger *et al.* (2004) used the square root of daily precipitation as a predictor, but we found that taking the natural logarithm of precipitation achieves a more symmetric distribution. To address the binary nature of the data, dry days must be omitted from the regression. However, simulated and observed time series of precipitation rarely contain the same amount of precipitable days. To correct for any differences in number of rainy days between observations and the simulated time series, the two time series are ordered by rank, extracting the top number of values in each vector corresponding to the number of rainy days in the shorter non-zero time series (usually observations, because AOGCMs tend to 'drizzle' or simulate many more low-precipitation days than observed; e.g. Chen *et al.*, 1996; Sun *et al.*, 2006; Perkins *et al.*, 2007). Drizzle is also addressed by setting downscaled precipitation amounts less than trace (typically defined as 0.005 inches or 0.127 mm) to zero.

The fact that the downscaling process can only be applied to precipitable days raises concerns regarding model performance in extremely dry regions. Given the typical variance of precipitation, to have a confidence level of 95% there must be at least 57 samples in the dataset (i.e. at least 57 wet days in each of the 12 monthly time series that span the entire training period). This value was determined by applying a simple sample size calculation for linear and logistic regression following

Hsieh *et al.* (1998). During the dry season in arid regions the sample size can be insufficient, even for 50 years of data including half a month on either side. If the sample is insufficient, the model automatically expands it by including an extra week's data on either side of that month (thus containing 3 weeks each of the prior and subsequent months), repeating the process up to a maximum of eight times until a sample size of at least 57 is reached. If 16 weeks have been added and the sample is still less than 57 but greater than 20, a linear regression is used. If less than 20 (which, for a training period of 50 years, would mean less than 1 day in 2 years with measurable precipitation), all downscaled values are set to zero for that month. This procedure has been tested and produces reasonable downscaling of historical precipitation in regions that are arid or semiarid.

The ARRM model was constructed in three distinct phases to quantify the contribution of specific elements to model performance. All phases build monthly models that incorporate 2 weeks' data on either side of the target month to double the sample size, and all versions prefilter the temperature and precipitation predictors as described above before ranking by value. The difference between the versions is the function used to fit the quantile–quantile relationship between observations and historical simulations. The first version applies a least mean squares linear fit (using the function *lm* in R), similar to that used in the SAR approach of Dettinger *et al.* (2004). The second version applies the piecewise regression function described above. The third version also uses piecewise regression, but incorporates removal of negative or flat slopes and bias correction near the tails. Removal of negative slopes is not expected to yield significant improvements in model performance, and in some cases it may even degrade initial performance; however, it is necessary to reduce the risk of unrealistic statistical relationships between modelled and observed values. The purpose of the comparison is not primarily to demonstrate the superiority of the final model, but rather to ensure that model performance is not overly degraded by this step.

The three different versions will be referred to as *linear*, *simple piecewise*, and *full piecewise* downscaling models, respectively. The ability of these three versions to downscale daily temperature and precipitation for 20 long-term stations in North America was evaluated using the data and model simulations described next.

3. Data and simulations

3.1. Observations

Downscaling was conducted and tested using observed daily minimum and maximum temperature and 24-h cumulative precipitation amounts for 20 long-term North American weather stations for the period 1960–2009. Seventeen of the stations are distributed across diverse climatic regions in the continental United States including coastal, central, and mountainous regions; two stations

are located in Canada; and one in Mexico (Figure 4). Records were obtained from the GHCN (Peterson and Vose, 1997).

Although GHCN station data have already undergone a standardized quality control (Durre *et al.*, 2008), before using the station data for downscaling they were filtered using a quality control algorithm to identify and remove (replacing with 'NA') erroneous values in the GHCN database. This additional quality control step included three tests for errors, removing data on any days where the daily reported minimum temperature exceeds the reported maximum, any temperature values above (below) the highest (lowest) recorded values for North America (-50 to 70°C) or with precipitation below zero or above the highest recorded value for the continental United States (915 mm in 24 h), and repeated values of more than five consecutive days with identical temperature or non-zero precipitation values to the first decimal. Additionally, an erroneous value was found for Hialeah, FL, of -17.8°C on 8 November 2003 (with temperatures of 29.4°C the previous day and 30.0°C the following day), which was removed.

3.2. Atmosphere–ocean general circulation models

Model output from four AOGCMs was used to evaluate the downscaling model. The models chosen for this study are all part of the Coupled Model Inter-comparison Project version 3 (Meehl *et al.*, 2007): the National Center for Atmospheric Research Community Climate System Model version 3 (CCSM3; Collins *et al.*, 2006), the National Oceanic and Atmospheric Administration/Geophysical Fluid Dynamics Laboratory Climate Model version 2.1 (GFDL CM2.1; Delworth *et al.*, 2006), the United Kingdom Met Office Climate Model version 3 (HadCM3; Pope *et al.*, 2000), and the Department Of Energy/National Center for Atmospheric Research Parallel Climate Model (PCM; Washington *et al.*, 2000). Previous studies (e.g. Gleckler *et al.*, 2008; Stoner *et al.*, 2009; Rusticucci *et al.*, 2010) show that these models are able to represent key features of atmospheric variability including teleconnection patterns, extreme temperature and precipitation, as well as other climate metrics. A description of the models is provided in Table I.

3.3. Scenarios

Historical AOGCM simulations correspond to the CMIP '20th Century Climate in Coupled Models' or 20C3M total forcing scenarios. These scenarios include forcing from anthropogenic emissions of greenhouse gases, aerosols, and reactive species; changes in solar output; particulate emissions from volcanic eruptions; changes in tropospheric and stratospheric ozone; and other influences required to provide a complete picture of the climate over the last century. Where multiple simulations were available, the first was used here (run 1).

To represent a broad range of alternative climate futures, simulations corresponding to the IPCC Special

Report on Emission Scenarios (SRES) higher (A1fi) and lower (B1) emission scenarios were used (Nakićenović *et al.*, 2000). These scenarios describe internally consistent pathways of future societal development and corresponding emissions, with atmospheric CO_2 concentrations reaching approximately 550 ppm (B1) and 990 ppm (A1fi) by 2100.

20C3M simulations only cover the period 1960–1999, in order to have a longer range of historical simulations we extended this period by 10 years by including 2000–2009 simulated output from the A2 SRES scenario. We find this to be a reasonable approach because the inertia of the climate system delays its response to forcings from increased greenhouse gases and other factors identifying each scenario and there is not much difference between the scenarios over the first decade of the century (Stott and Kettleborough, 2002).

4. Model evaluation

4.1. Creation of independent simulated historical time series

To evaluate the three versions of ARRM (linear, simple piecewise, and full piecewise), 50 years' worth of data and historical total forcing simulations from 1960 to 2009 were used to build downscaling models for daily temperature and precipitation for 20 long-term weather stations across North America. *N*-fold cross-validation, or jackknifing, was used whereby the downscaling model was trained on all but one of the years, then used to predict values for the remaining year. ARRM builds a separate model for each of the 12 months of the year, so this process was repeated until 600 independent 1-month simulated daily time series had been generated for each location, independent of the observations used to train the statistical model. These were then combined into a single 50-year time series for evaluation.

Use of cross-validation in creating the historical simulated time series to be evaluated against observations is a crucial aspect of the evaluation. If the statistical model had been trained on all 50 years and then used to predict those same 50 years, comparing the resulting time series with observations would simply reflect the ability of the regression function to fit the data. The results of such an evaluation would be *improved* by overfitting, for example, by allowing the piecewise regression function to fit an infinite number of knot points to the quantile–quantile relationship. In contrast, by generating an independent time series, the evaluation instead reflects the ability of the model to recreate observations that were not used to train the model. The results of such an evaluation are *degraded* by overfitting that makes the model less generalizable. The split-sample approach, whereby observational data are divided into a training and evaluation period, is commonly used to evaluate statistical downscaling methods in the literature. The ability of the statistical model to reproduce

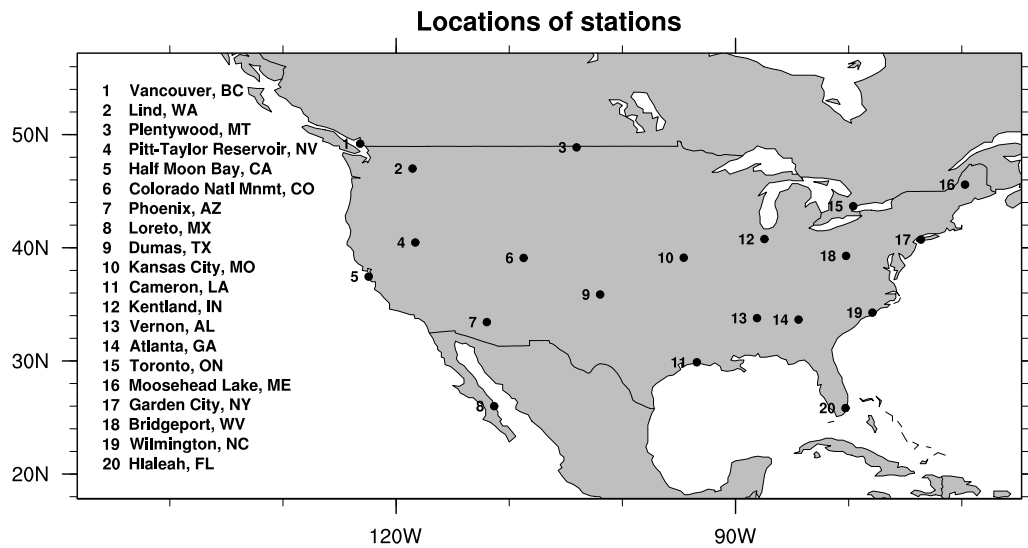
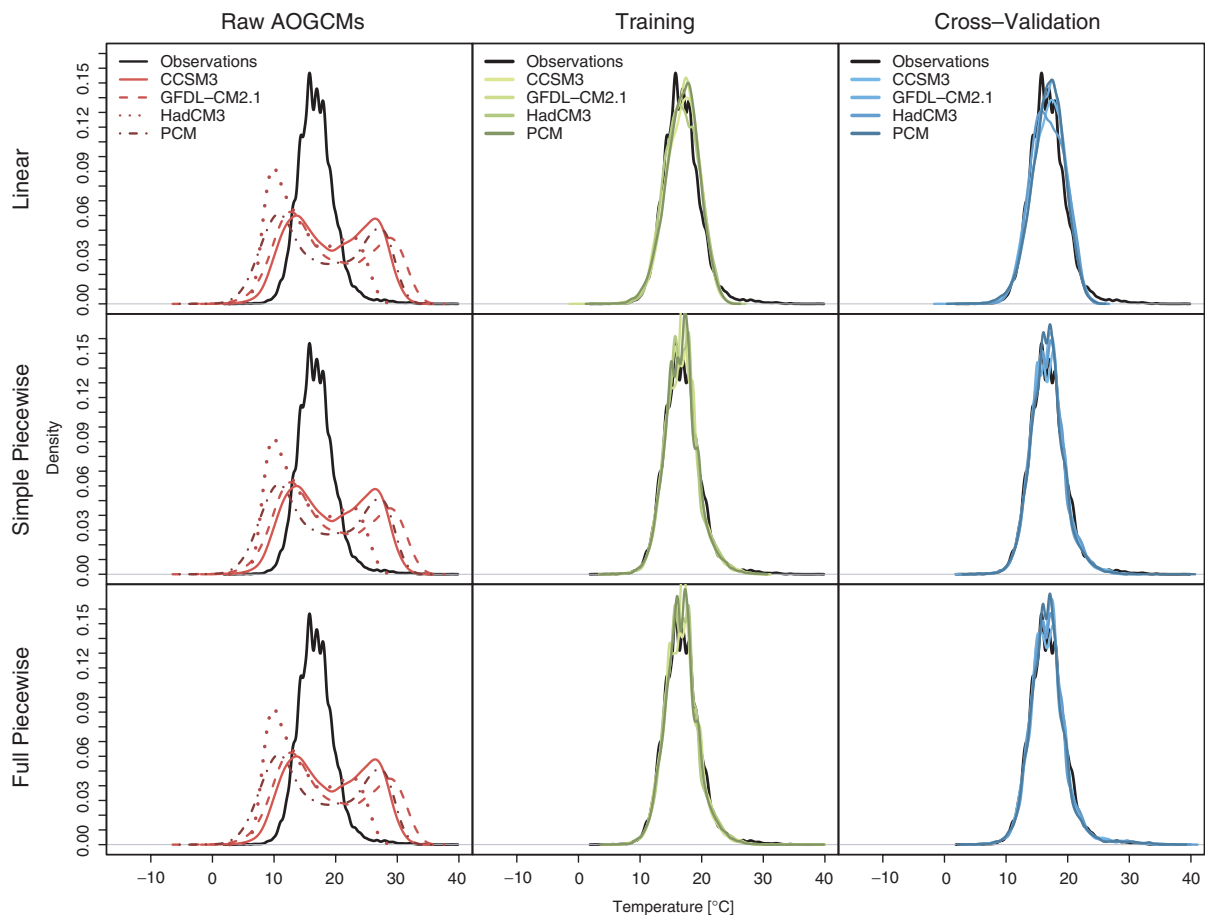


Figure 4. Locations of the 20 North American stations used to evaluate ARRMs.

Figure 5. Probability distributions for Half Moon Bay, CA, 1960–2009 observed daily maximum temperature (black) and AOGCM (left column), training (middle column), and cross-validation (right column) simulated daily maximum temperature using linear (top row), simple piecewise (middle row), and full piecewise (bottom row) regression models. This figure is available in colour online at wileyonlinelibrary.com/journal/joc

observed natural variability at a given location, however, depends on the degree to which it is able to sample from that variability in both training and evaluation. The split-sample approach limits the sample size of both the training and observation periods (typically, $N/2$ years

each), whereas the jackknifed cross-validation approach used here, with a training period of $N - 1$ years and an evaluation period of N years, more closely approximates the skill of the full dependent sample model that will be used to downscale future projections. As the purpose

Table I. Summary of key characteristics of AOGCMs used, including acronyms, host institution, as well as atmospheric and oceanic resolution.

Model acronym	Host institution	Atmospheric resolution	Oceanic resolution
CCSM3	NCAR (USA)	$1.4^{\circ} \times 1.4^{\circ}$	$1.125^{\circ} \times 0.43^{\circ}$
GFDL-CM2.1	NOAA/GFDL (USA)	$2.0^{\circ} \times 2.5^{\circ}$	$0.9^{\circ} \times 1.0^{\circ}$
HadCM3	UK Met Office (UK)	$2.5^{\circ} \times 3.75^{\circ}$	$1.25^{\circ} \times 1.25^{\circ}$
PCM	NCAR (USA)	$2.81^{\circ} \times 2.81^{\circ}$	$1.0^{\circ} \times 1.0^{\circ}$

of downscaling is to ‘recreate’ future conditions that cannot be used to train the model, we argue that the type of evaluation done here is more relevant to assessing the performance of a downscaling model. This is somewhat similar to a bootstrapping approach (Li *et al.*, 2010).

4.2. Evaluating temperature downscaling

The overall skill of the downscaling models is assessed in terms of their ability to reproduce the observed annual distribution (through comparing probability distribution functions), the RMSEs compared to observations, and the absolute value of the bias in the 0.1th, 1st, 10th, 25th, 50th, 75th, 90th, 99th, and 99.9th quantiles. Model projections are also compared (although not evaluated) for end-of-century under the SRES A1fi (higher) and B1 (lower) emissions scenarios.

To gain a qualitative perspective on the downscaling, we first compare observed, AOGCM-simulated (nearest grid cell), downscaled (training period), and downscaled (independent evaluation period) maximum and minimum temperature distributions for the coastal location of Half Moon Bay, CA (Figures 5 and 6), for which the simulated and observed temperature distributions differ noticeably. The three rows correspond to the three versions of the downscaling model (linear, simple piecewise, and full piecewise). The three columns show AOGCM predictions, predictions from training the downscaling model on all 50 years of data, and the independent cross-validation predictions, derived by the method described above. Identical figures for the remaining 19 stations and other graphics not included in this publication are available online (<http://temagami.ttu.edu/armm/>).

For this location, all AOGCMs simulate a wide distribution for maximum temperature with two peaks near 10 and 28 °C (Figure 5). In contrast, the distribution of observed maximum temperatures is narrow and only has one peak near 17 °C. The HadCM3 distribution is additionally skewed towards lower temperatures. One reason for the large difference between observed and simulated distributions is due to the landmask in the AOGCMs, which can have anything from 0 to 100% land fraction in coastal grid cells, differing between AOGCMs. Table II gives land fraction values for grid cells used to downscale stations near the coast. The grid cell downscaled to Half Moon Bay has only partial land coverage in most models (PCM: 15.2%, CCSM3: 53.8%, and GFDL-CM2.1: 84.2%) and is a complete ocean grid cell in the HadCM3 model. Predictions might be improved by selection of a

Table II. Fraction of land (in percent) of AOGCM grid cell used to downscale each station for the four AOGCMs. Values are given only for stations near a coast as the percentage of land in grid cells used to downscale inland stations were all 100%.

Station	CCSM3	GFDL-CM2.1	HadCM3	PCM
Cameron, LA	85	66	100	42
Garden City, NY	91	84	100	44
Half Moon Bay, CA	54	84	0	15
Hialeah, FL	31	10	0	55
Loreto, MX	2	37	0	40
Moosehead Lake, ME	100	95	100	100
Phoenix, AZ	100	100	100	94
Vancouver, BC	100	100	100	88
Wilmington, NC	51	39	100	89

different AOGCM grid cell; however, the purpose here is not to generate optimal predictions but rather to test the ability of the downscaling method to correct AOGCM output. From that perspective, using a near-shore grid cell to simulate coastal conditions represents a greater challenge for the model, and all three versions of the downscaling model are able to approximate observations for these grid cells, narrowing the simulated distribution and removing the double peaks. The linear model is able to capture the general shape of the observed distribution, but underestimates high temperatures towards the tail of the distribution. This is improved upon by the simple piecewise model and almost completely resolved by the full piecewise model. There is little difference between the results for the training (middle column) and independent (last column) predictions, indicating that the downscaling model does not overfit and is successful at simulating observed conditions outside the training period.

There are some differences between maximum and minimum temperature (Figure 6). First, AOGCM distributions for minimum temperature more closely resemble observed, although generally skewed towards cooler when compared with warmer values. Second, all three downscaling models perform well at the tails of the distribution, but the peak of the distribution is better resolved by the two downscaling models that apply the piecewise regression technique.

Figure 7 compares the RMSE in maximum temperature across the entire distribution for all 20 stations. Applying any of the three downscaling models greatly reduces RMSEs compared with raw AOGCM outputs, which in most cases are an order of magnitude larger. Moreover, the downscaling process is able to transform a broad




Figure 6. Probability distributions for Half Moon Bay, CA, 1960–2009 observed daily minimum temperature (black) and AOGCM (left column), training (middle column), and cross-validation (right column) simulated daily minimum temperature using linear (top row), simple piecewise (middle row), and full piecewise (bottom row) regression models. This figure is available in colour online at wileyonlinelibrary.com/journal/joc

range of AOGCM predictions into distributions closely resembling observed. For all 20 stations, downscaling reduces the RMSE of simulated historical values from 2 to 8 °C down to less than 0.5 °C. Refining the downscaling technique by applying piecewise regression further decreases the residuals. There is little difference between RMSEs of the simple piecewise and full piecewise regression methods as improvements due to bias removal tend to be offset by removal of negative slopes. Results for minimum temperature (not shown) are similar, except that the RMSE values for AOGCMs tend to be lower, confirming the indication from Figures 5 and 6 that these AOGCMs are generally better at simulating daily minimum when compared with maximum temperatures, regardless of location.

The results of this evaluation are summarized by scatter plots of downscaled versus AOGCM RMSE (Figure 8). Applying downscaling reduces the spread of RMSEs noticeably with the linear version of the downscaling model, and even further when piecewise regression is added to the downscaling model, with RMSE values below 0.5 °C for temperature and below 10 mm for precipitation. For both simple piecewise and full piecewise downscaling models, the majority of points

are clustered between 0.2 and 0.3 °C for temperature and between 1 and 5 mm for precipitation (the far outlier for precipitation for the simple model is HadCM3 downscaled for Phoenix, AZ, with an RMSE value of 76.2 mm), indicating that this level of bias is most likely the limit to the ability of this particular type of statistical downscaling model, within the range of natural variability represented in the training dataset.

The third measure used to evaluate the downscaling methods is by examining the bias in the 0.1th, 1st, 10th, 25th, 50th, 75th, 90th, 99th, and 99.9th quantiles (Figure 9). Bias in AOGCM output is generally an order of magnitude larger than bias downscaled output, regardless of downscaling technique. There is no consistent tendency for AOGCM biases to be larger for certain quantiles, but downscaled quantiles tend to be slightly larger for extreme when compared with median quantiles.

Figure 9 also shows the percentage missing data in the observations for each station. Even for locations with a very high percentage of missing data (Bridgeport and Moosehead Lake) downscaling is able to improve on AOGCM output, although the resulting biases reflect the uncertainty from the very small sample size of the data used to train the statistical models.

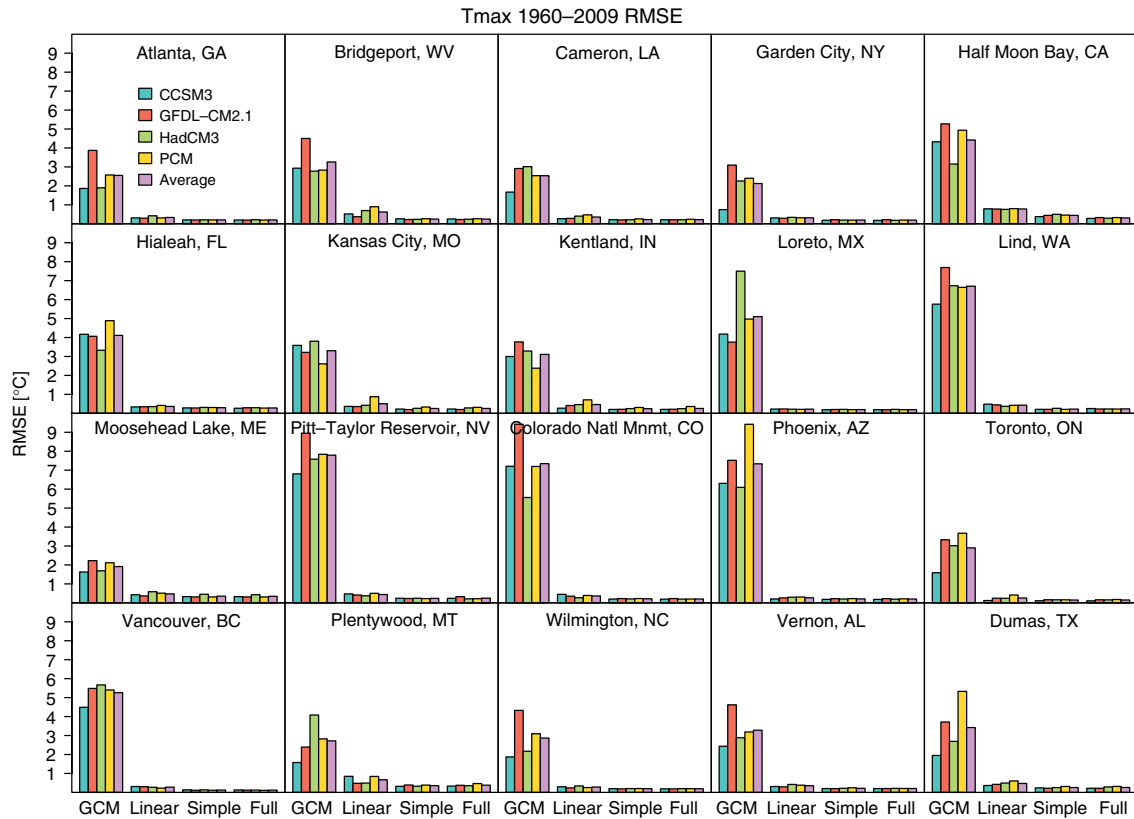


Figure 7. Daily maximum temperature root-mean-square errors, relative to observations, for raw AOGCM simulations as well as linear, simple piecewise, and full piecewise downsampling methods for each of the 20 stations. This figure is available in colour online at wileyonlinelibrary.com/journal/joc




Figure 8. Scatter plot of raw AOGCM 1960–2009 daily maximum, minimum temperature and precipitation RMSE *versus* cross-validation downscaled daily maximum temperature RMSE for all 20 stations for linear, simple piecewise, and full piecewise downsampling methods. Each dot denotes one station/downscaled AOGCM combination. This figure is available in colour online at wileyonlinelibrary.com/journal/joc




Figure 9. Absolute value of bias in nine quantiles between cross-validation downscaled 1960–2009 daily maximum temperature and observed daily maximum temperature for the same period. Values are averaged across the four AOGCMs and shown for raw, linear, simple piecewise, and full piecewise downscaling methods. The percentage shown in the upper right-hand corner of each plot is the percentage of missing data in the observed record. This figure is available in colour online at wileyonlinelibrary.com/journal/joc

Comparing the reduction in biases in the lowest, middle, and highest quantiles of maximum temperature achieved by downscaling from AOGCM outputs for the cross-validation results shows that using the linear downscaling method noticeably reduces the range in bias relative to AOGCM output for the median quantiles, but not for more extreme quantiles (Figure 10; results for minimum temperature are similar, not shown). Incorporating piecewise regression makes little difference to the 50th quantile when compared with the linear model, but significantly reduces biases in more extreme quantiles.

This suggests that the piecewise regression technique's primary improvement for temperature, compared with a linear model, is in downscaling extreme values.

4.3. Evaluating precipitation downscaling

To gain a qualitative perspective on precipitation downscaling, we first compare observed, AOGCM-based, and downscaled distributions of the natural logarithm of precipitation for 1960–2009 for Kentland, IN (Figure 11). The left column of plots shows the tendency of AOGCMs to drizzle on the left-hand side of the distribution




Figure 10. Scatter plots of raw AOGCM maximum temperature bias *versus* downscaled bias for the evaluation period for three separate quantiles representing the tail, the middle, and high ends of the distribution. This relationship is shown for linear (top row), simple piecewise (middle row), and full piecewise (bottom row) downscaling methods. This figure is available in colour online at wileyonlinelibrary.com/journal/joc

and underestimates the magnitude of high-precipitation extremes on the right-hand side of the distribution (e.g. 150 vs 400 mm). The AOGCMs also fail to simulate the double-peaked distribution common to many stations, including Kentland.

The linear version of the downscaling model corrects the lower tail, partly corrects the higher tail (although it introduces some very high-precipitation values), and does not correct for the two peaks in the distribution. Incorporating piecewise regression resolves the peaks but introduces artificially large extreme values that are corrected in the full piecewise method that includes bias correction.

Figure 12 compares RMSE values for all 20 stations between observations, AOGCM output, and downscaled simulations for the evaluation period. For almost all locations, applying the linear downscaling model increases RMSEs relative to AOGCM output. This is most likely due to the linear model simulating extreme values that are too high but carry more weight in the overall RMSE calculations. Piecewise regression corrects the high-end bias and in almost all cases reduces RMSE relative to AOGCM output.

Absolute bias (in percent) in the same nine quantiles as used for temperature (Figure 13) shows that for all nine quantiles, biases are generally small, for the full model

the bars are barely visible for most stations for all nine quantiles. Plotting real-value quantile biases for the 0.1th, 50th, and 99.9th quantiles (Figure 14) shows again that biases are very minimal for the lower and middle quantiles, with larger values for the highest quantile. AOGCM biases in the 99.9th quantile are nearly all negative, i.e. AOGCMs underestimate extreme precipitation accumulation in almost all 20 locations examined here. This is not surprising, given that AOGCM values are averaged over a large area whereas observations are for point sources.

5. Future projections

The purpose of most downscaling models is to generate future projections more representative of individual locations than current AOGCMs can provide with grid cell-sized information. Here, we compare the results of AOGCM simulations with ARRM downscaled future projections using the entire historical period (1960–2009) to train each model.

5.1. Maximum temperature

Figures 17 and 18 show the change in downscaled *versus* raw AOGCM daily maximum temperature for 2070–2099 relative to the historical period observations




Figure 11. Probability distributions for 1960–2009 observed daily precipitation accumulation (black lines), AOGCM simulated daily precipitation accumulation (left column), training (middle column), and cross-validation (right column) for Kentland, IN. Downscaled simulations are shown for linear (top row), simple piecewise (middle row), and full piecewise (bottom row) methods. This figure is available in colour online at wileyonlinelibrary.com/journal/joc




Figure 12. Daily precipitation root-mean-square errors, relative to observations, for raw AOGCM simulations as well as linear, simple piecewise, and full piecewise downscaling methods for the evaluation period for each of the 20 stations. This figure is available in colour online at wileyonlinelibrary.com/journal/joc




Figure 13. Percentage bias in nine quantiles between cross-validation downscaled and observed 1960–2009 daily precipitation accumulation. Values are averaged across the four AOGCMs and shown for raw, linear, simple piecewise, and full piecewise downscaling methods. The percentage shown in the upper right-hand corner of each plot is the percentage of missing data in the observed record. This figure is available in colour online at wileyonlinelibrary.com/journal/joc

(1960–2009) for the three temperature downscaling models (linear, simple piecewise, and full piecewise) and 0.1th, 50th, and 99.9th quantiles. Under the higher A1fi scenario (Figure 15), the most obvious difference between raw AOGCM *versus* downscaled future changes is that downscaling produces only positive changes (i.e. increases) in all three quantiles illustrated [with the exception of one station (Half Moon Bay, CA) for the linear model and 99.9th quantile], whereas raw AOGCM changes are both positive as well as negative for these

three quantiles, indicating that the raw output projects warming for some locations and cooling for others at the end of the century. For the lower B1 scenario (Figure 16), more stations also show warming at the end of the century after downscaling compared with raw AOGCM results, especially for the middle and upper quantiles. Some cooling is projected for the lowest quantiles, indicating that some stations might see a wider distribution of daily maximum temperature at the end of the century with more extremes in both ends of the distribution.




Figure 14. Scatter plots of raw AOGCM quantile bias *versus* downscaled quantile bias for three separate quantiles representing the low tail, the middle, and high tail of the distribution. This relationship is shown for linear, simple piecewise, and full piecewise downscaling methods for the evaluation period. This figure is available in colour online at wileyonlinelibrary.com/journal/joc




Figure 15. Scatter plots of raw A1fi AOGCM maximum temperature (2070–2099)–(1960–2009) change *versus* downscaled change for three separate quantiles representing the tail, the middle, and the high ends of the distribution. This relationship is shown for linear (top), simple piecewise (middle), and full piecewise (bottom) downscaling methods. This figure is available in colour online at wileyonlinelibrary.com/journal/joc




Figure 16. Scatter plots of raw B1 AOGCM maximum temperature (2070–2099)–(1960–2009) change *versus* downscaled change for three separate quantiles representing the tail, the middle, and high ends of the distribution. This relationship is shown for linear (top), simple piecewise (middle), and full piecewise (bottom) downscaling methods. This figure is available in colour online at wileyonlinelibrary.com/journal/joc

Figure 17 shows the mean AOGCM absolute 2070–2099 daily maximum temperature changes, relative to 1960–2009, in each of the nine quantiles for the A1fi and B1 scenarios for raw (light-coloured bars) and downscaled output (dark-coloured bars). There is overall a general agreement among the 20 stations that a greater change in the 50th quantile is projected for the A1fi scenario than for the B1 scenario (Loreto, MX, and Hialeah, FL, being the only exceptions – note that these stations also have low land fraction in all four AOGCMs). However, there is no general tendency for the mean change to be more or less for downscaled *versus* raw AOGCM output, with some locations, such as Atlanta, GA, and Bridgeport, WV, showing a larger change projected for the A1fi scenario than for the B1 scenario (Loreto, MX, and Hialeah, FL, being the only exceptions). Similarly, projected changes in higher quantiles from raw AOGCM can be higher than downscaled for certain locations and lower for others. This indicates that downscaling produces results specific to each location, as opposed to the more general AOGCM grid cell output.

5.2. Precipitation

Figure 18 shows the 2070–2099 relative to 1960–2009 raw AOGCM *versus* downscaled precipitation changes in the 0.1th, 50th, and 99.9th quantiles for the three versions

of the downscaling model for the A1fi scenario, given as RMSE differences. Unsurprisingly, there is less than 1 mm change in predicted changes for the 0.1th quantile for all 20 stations for both raw AOGCM and downscaled projections. The reason for the fixed RMSE values for the 0.1th quantile for the raw AOGCMs is due to weather stations not reporting *trace* precipitation, which is set at 0.005 inches (0.127 mm). The higher frequency of low precipitation events, compared with higher precipitation events, in most locations causes the 0.1th quantile to almost always equal to the lowest recorded or simulated precipitation value. The lowest simulated value in AOGCMs, when rounded to the nearest 2 decimals, is 0.01 mm, because AOGCMs do not allow for ‘trace’. All but one stations have a lowest value of 0.2 mm (when converted from inches), whereas one station (Bridgeport, WV) has 0.1 mm as the lowest recorded value, which is the cause for that station not being in agreement with the others in the bias plot (Figure 18). Under both scenarios, AOGCM outputs project little to some (up to about 8 mm) decrease in the amount of precipitation comprising the median quantile, whereas when downscaled the same quantile shows less than 2 mm change from current conditions, with few exceptions, for both scenarios. The largest change is in the 99.9th quantile for both scenarios, with up to several hundred millimetres change

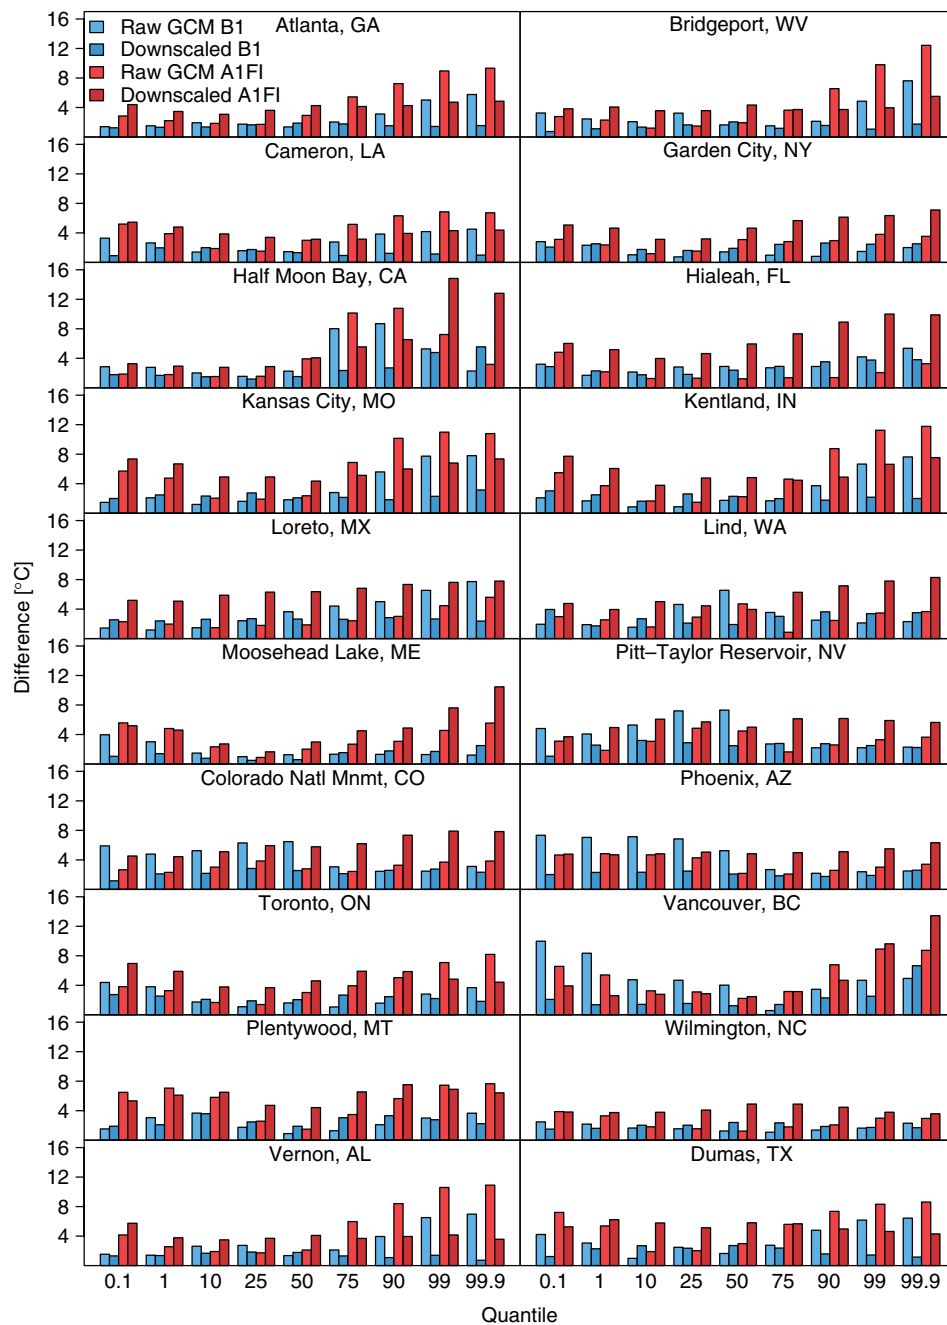


Figure 17. Absolute change in nine quantiles between 2070–2099 simulated and 1960–2009 observed daily maximum temperature. Values are averaged across the four AOGCMs and shown for raw and downscaled (full piecewise) A1fi and B1 scenarios. This figure is available in colour online at wileyonlinelibrary.com/journal/joc

from current extreme conditions. For raw AOGCM projections, future extreme precipitation amounts appear to have decreased, whereas when downscaled, the same locations show a large increase in precipitation extremes, especially for the A1fi scenario (Figure 18). This is most likely due to poor simulation of precipitation at the local scale by AOGCMs and is corrected by applying the statistical downscaling model, which is trained on historical temporal precipitation variability for each location. The linear version of the downscaling model produces very large, up to about 1750 mm, increases in extreme precipitation events, whereas the full piecewise

downscaling model produces more moderate, but still large – up to about 300 mm increases in extreme events. The numbers are very similar, although slightly smaller, for the B1 scenario (not shown here, but available at <http://temagami.ttu.edu/arm/>).

Absolute precipitation changes for the 99th and 99.9th quantiles are shown in Figure 19 for both A1FI and B1 scenarios, averaged across all four AOGCMs. For some stations, such as Hialeah, FL, Loreto, MX, and Vernon, AL, there is a substantial difference between extreme event projections for raw *versus* downscaled AOGCMs, with the raw AOGCM generally projecting




Figure 18. Scatter plots of raw AOGCM A1fi quantile difference between 2070–2099 simulated and 1960–2009 observed daily precipitation *versus* downscaled quantile difference for three separate quantiles representing the low tail, the middle, and high tail of the distribution. This relationship is shown for linear, simple piecewise, and full piecewise downsampling methods. This figure is available in colour online at wileyonlinelibrary.com/journal/joc

larger absolute changes, compared with present conditions, than downscaled projections. The results of the cross-validation evaluation suggest that more confidence could be placed in the downscaled projections when compared with raw AOGCM output, because downscaled projections are tailored to each individual location.

6. Conclusions

The ARRM is an empirical statistical downscaling model capable of downscaling local projections of temperature and precipitation to both station-based observations and spatially gridded observations. Quantile regression, the method on which ARRM is based, is unique in that it builds a regression model based on matching the quantiles of the observed and simulated time series as opposed to matching corresponding day-to-day data points, which is the basis for many other regression-based statistical downscaling studies (Wilby *et al.*, 1998; Huth, 1999, 2002; Wilby and Wigley, 2000; Huth *et al.*, 2001; Boé *et al.*, 2007; Kostopoulou *et al.*, 2007). ARRM adds to this by using a piecewise regression model instead of a straight linear regression, which improves its ability to simulate more extreme temperatures and precipitation, one of the major issues with other downscaling methods (Huth, 1999; Goodess *et al.*, 2012).

The downscaling model was evaluated based on cross-validation of three different (linear, simple piecewise, and full piecewise) versions of both the temperature and precipitation models. Each version was evaluated in terms of three different quantities: the distributions, giving a visual image of the skill each model; the RMSE; and bias in a range of quantiles.

The addition of piecewise regression, instead of straight linear regression, was found to have the largest impact on the performance of the method. The largest biases were found to be near the tails of the distribution, primarily due to data sparseness. Some sensitivity to station location was found in the linear versions of the downscaling model, but the addition of piecewise regression was able to eliminate much of this.

For future projections, the spread among projected temperature increases is generally narrower for downscaled temperature compared with raw AOGCM projections for the three quantiles shown, with more stations showing positive temperature changes after downscaling than before, for both higher A1fi and lower B1 scenarios. Downscaled projections of precipitation show smaller changes for the 50th quantile than raw AOGCM projections, for both A1fi and B1 scenarios, but slightly larger changes in extreme events, with projected changes being generally greater under the A1fi scenario than the B1 scenario.




Figure 19. Absolute bias in the 99th and 99.9th quantiles between 2070–2099 simulated and 1960–2009 observed daily precipitation accumulation. Values are averaged across the four AOGCMs. This figure is available in colour online at wileyonlinelibrary.com/journal/joc

Evaluating the ability of ARRM to reproduce observed temperature and precipitation at 20 stations across North America shows that the statistical downscaling model is able to reproduce values from the 0.1th to the 99.9th quantiles with biases generally below 1°C and 5 mm. Downscaling future projections can alter the sign of AOGCM-simulated changes and usually narrows the range of projected changes across multiple AOGCM simulations.

The ultimate purpose of the ARRM framework is to allow for user selection from a broad range of predictors and predictands to efficiently downscale either point source or gridded observations of any observed

climate variable with a Gaussian-like distribution that can be predicted from large-scale AOGCM output fields. Model performance for station-based temperature and precipitation downscaling appears sufficient to support continued development of such a generalized model. Future work will describe the application of this model to gridded datasets and to downscaling solar radiation and relative humidity.

Acknowledgements

This work was supported by the National Science Foundation, under NSF contract number DMS-0724377,

and the U.S. Geological Survey, under contract number G10AC00248. Station data were obtained from the Global Historical Climatology Network, and AOGCM output fields were retrieved from the Earth System Grid.

References

- Boé J, Terray L, Habets F, Martin E. 2007. Statistical and dynamical downscaling of the Seine basin climate for hydro-meteorological studies. *International Journal of Climatology* **27**: 1643–1655. DOI: 10.1002/joc.1602.
- Chen M, Dickinson RE, Zeng X, Hahmann AN. 1996. Comparison of precipitation observed over the continental United States to that simulated by a climate model. *Journal of Climate* **9**: 2233–2249. DOI: 10.1175/1520-0442(1996)009<2233:COPOOT>2.0.CO;2.
- Collins WD, Bitz CM, Blackmon ML, Bonan GB, Bretherton CS, Carton JA, Chang P, Doney SC, Hack JJ, Henderson TB, Kiehl JT, Large WG, McKenna DS, Santer BD, Smith RD. 2006. The Community Climate System Model Version 3 (CCSM3). *Journal of Climate* **19**: 2122–2143. DOI: 10.1175/JCLI3761.1.
- Crane RG, Hewitson BC. 1998. Doubled CO₂ precipitation changes for the Susquehanna basin: down-scaling from the GENESIS general circulation model. *International Journal of Climatology* **18**: 65–76. DOI: 10.1002/(SICI)1097-0088(199801)18:1<65::AID-JOC222>3.0.CO;2-9.
- Delworth TL, Broccoli AJ, Rosati A, Stouffer RJ, Balaji V, Beesley JA, Cooke WF, Dixon KW, Dunne J, Dunne KA, Durachta JW, Findell KL, Ginoux P, Gnanadesikan A, Gordon CT, Griffies SM, Gudgel R, Harrison MJ, Held IM, Hemler RS, Horowitz LW, Klein SA, Knutson TR, Kushner PJ, Langenhorst AR, Lee HC, Lin SJ, Lu J, Malyshev SL, Milly PCD, Ramaswamy V, Russell J, Schwarzkopf MD, Shevliakova E, Sirutis JJ, Wyman B, Zeng F, Zhang R. 2006. GFDL's CM2 global coupled climate models. Part I: Formulation and simulation characteristics. *Journal of Climate – Special Section* **19**: 643–674. DOI: 10.1175/JCLI3629.1.
- Dettinger MD, Cayan DR, Meyer MK, Jeton AE. 2004. Simulated hydrologic responses to climate variations and change in the Merced, Carson, and American river basins, Sierra Nevada, California, 1900–2099. *Climatic Change* **62**: 283–317.
- Durre I, Vose RS, Wuertz DB. 2008. Robust automated quality assurance of radiosonde temperatures. *Journal of Applied Meteorology and Climatology* **47**: 2081–2095.
- Gleckler PJ, Taylor KE, Doutriaux C. 2008. Performance metrics for climate models. *Journal of Geophysical Research* **113**: D06104. DOI: 10.1029/2007JD008972.
- Goodess CM, Anagnostopoulou C, Bárdossy A, Frei C, Harpham C, Haylock MR, Hundecha Y, Maheras P, Ribalaygua J, Schmidli J, Schmith T, Tolika K, Tomozeiu R, Wilby RL. 2012. An intercomparison of statistical downscaling methods for Europe and European regions – assessing their performance with respect to extreme temperature and precipitation events. University of East Anglia: UK, Climate Research Unit Research Publications. Report No. CRU RP11:1–72.
- Hay LE, Wilby RL, Leavesley GH. 2000. A comparison of delta change and downscaled GCM scenarios for three mountainous basins in the United States. *Journal of the American Water Resources Association* **36**: 387–397. DOI: 10.1111/j.1752-1688.2000.tb04276.x.
- Hayhoe K, Cayan D, Field CB, Frumhoff PC, Maurer EP, Miller NL, Moser SC, Schneider SH, Nicholas Cahill K, Cleland EE, Dale L, Drapek R, Hanemann RM, Kalkstein LS, Lenihan J, Lunch CK, Neilson RP, Sheridan SC, Verville JH. 2004. Emissions pathways, climate change, and impacts on California. *Proceedings of the National Academy of Sciences* **101**: 12422–12427.
- Hayhoe K, Wake C, Anderson B, Liang XZ, Maurer E, Zhu J, Bradbury J, DeGaetano A, Stoner AM, Wuebbles D. 2008. Regional climate change projections for the Northeast USA. *Mitigation and Adaptation Strategies for Global Change* **13**: 425–436. DOI: 10.1007/s11027-007-9133-2.
- Hayhoe K, VanDorn J, Croley TC II, Schlegal N, Wuebbles D. 2010. Regional climate change projections for Chicago and the US Great Lakes. *Journal of Great Lakes Research* **36**: 7–21. DOI: 10.1016/j.jglr.2010.03.012.
- Haylock MR, Cawley GC, Harpham C, Wilby RL, Goodess CM. 2006. Downscaling heavy precipitation over the United Kingdom: a comparison of dynamical and statistical methods and their future scenarios. *International Journal of Climatology* **26**: 1397–1415. DOI: 10.1002/joc.1318.
- Hidalgo HG, Dettinger MD, Cayan DR. 2008. Downscaling with constructed analogues: daily precipitation and temperature fields over the United States. California Energy Commission: California. Report No. CEC-500-2007-123: 1–62.
- Hsieh FY, Bloch DA, Larsen MD. 1998. A simple method of sample size calculation for linear and logistic regression. *Statistics in Medicine* **17**: 1623–1634.
- Huth R. 1999. Statistical downscaling in central Europe: evaluation of methods and potential predictors. *Climate Research* **13**: 91–101.
- Huth R. 2002. Statistical downscaling of daily temperature in central Europe. *Journal of Climate* **15**: 1732–1742. DOI: 10.1175/1520-0442(2002)015<1731:SDODTI>2.0.CO;2.
- Huth R, Kysely J, Dubrovsky M. 2001. Time structure of observed, GCM-simulated, downscaled and stochastically generated daily temperature series. *Journal of Climate* **6**: 4047–4061. DOI: 10.1175/1520-0442(2001)014<4047:TSOOGS>2.0.CO;2.
- Jeong DI, St-Hilaire A, Ouarda TBMJ, Gachon P. 2012. CGCM3 predictors used for daily temperature and precipitation downscaling in Southern Quebec, Canada. *Theoretical and Applied Climatology* **107**: 389–406. DOI: 10.1007/s00704-011-0490-0.
- Knutti R, Furrer R, Tebaldi C, Cernak J, Meehl GA. 2010. Challenges in combining projections from multiple climate models. *Journal of Climate* **23**: 2739–2758. DOI: 10.1175/2009JCLI3361.1.
- Koenker R, Bassett G Jr. 1978. Regression quantiles. *Econometrica* **46**: 33–50.
- Koenker R, Hallock KF. 2001. Quantile regression. *Journal of Economic Perspectives* **15**: 143–156.
- Kostopoulou E, Giannakopoulos C, Anagnostopoulou C, Tolika K, Maheras P, Vafiadis M, Founda D. 2007. Simulating maximum and minimum temperature over Greece: a comparison of three downscaling techniques. *Theoretical and Applied Climatology* **90**: 65–82. DOI: 10.1007/s00704-006-0269-x.
- Li H, Sheffield J, Wood EF. 2010. Bias correction of monthly precipitation and temperature fields from Intergovernmental Panel on Climate Change AR4 models using equidistant quantile matching. *Journal of Geophysical Research* **115**: D10101. DOI: 10.1029/2009JD012882.
- Luers AL, Cayan DR, Franco G, Hanemann M, Croes B. 2006. Our changing climate, assessing the risks to California. California Energy Commission: California. Report no. CEC-500-2006-077: 1–16.
- Mearns LO, Gutowski WJ, Jones R, Leung LY, McGinnis S, Nunes AMB, Qian Y. 2009. A regional climate change assessment program for North America. *Eos* **90**: 311–312. DOI: 10.1029/2009EO360002.
- Meehl GA, Covey C, Delworth T, Latif M, McAvaney B, Mitchell JFB, Stouffer RJ, Taylor KE. 2007. The WCRP CMIP3 multi-model dataset: a new era in climate change research. *Bulletin of the American Meteorological Society* **88**: 1383–1394. DOI: 10.1175/BAMS-88-9-1383.
- Nakićenović N, Alcamo J, Davis G, de Vries B, Fenhann J, Gaffin S, Gregory K, Grübler A, Jung TY, Kram T, La Rovere EL, Michaelis L, Mori S, Morita T, Pepper W, Pitcher H, Price L, Riahi K, Roehrl A, Rogner HH, Sankovski A, Schlesinger M, Shukla P, Smith S, Swart R, van Rooijen S, Victor N, Dadi Z. 2000. Special Report on Emission Scenarios. *Intergovernmental Panel on Climate Change*, Retrieved 1 April 2011. Available at http://www.grida.no/publications/other/ipcc_sr/
- O'Brien TP, Sornette D, McPherron RL. 2001. Statistical asynchronous regression: determining the relationship between two quantities that are not measured simultaneously. *Journal of Geophysical Research* **106**: 13247–13259. DOI: 10.1029/2000JA900193.
- Perkins SE, Pitman AJ, Holbrook NJ, McAneney J. 2007. Evaluation of the AR4 climate models' simulated daily maximum temperature, minimum temperature and precipitation over Australia using probability density functions. *Journal of Climate* **20**: 4356–4376. DOI: 10.1175/JCLI4253.1.
- Peterson TC, Vose RS. 1997. An overview of the Global Historical Climatology Network temperature database. *Bulletin of the American Meteorological Society* **78**: 2837–2849. DOI: 10.1175/1520-0477(1997)078<2837:AOTGH>2.0.CO;2.
- Pope VD, Gallani ML, Rowntree PR, Stratton RA. 2000. The impact of new physical parametrizations in the Hadley Centre climate model: HadAM3. *Climate Dynamics* **16**: 123–146. DOI: 10.1007/s003820050009.
- R Development Core Team. 2012. *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing: Vienna, Austria. ISBN 3-900051-07-0. Available at <http://www.R-project.org/>

- Raisanen J, Ylhäisi JS. 2011. How much should climate model output be smoothed in space? *Journal of Climate* **24**: 867–880. DOI: 10.1175/2010JCLI3872.1.
- Rusticucci M, Marengo J, Penalba O, Renom M. 2010. An intercomparison of model-simulated in extreme rainfall and temperature events during the last half of the twentieth century. Part 1: Mean values and variability. *Climatic Change* **98**: 493–508. DOI: 10.1007/s10584-009-9742-8.
- Schmidli J, Frei C, Vidale PL. 2006. Downscaling from GCM precipitation: a benchmark for dynamical and statistical downscaling methods. *International Journal of Climatology* **26**: 679–689. DOI: 10.1002/joc.1287.
- Stehlik J, Bardossy A. 2002. Multivariate stochastic downscaling model for generating daily precipitation series based on atmospheric circulation. *Journal of Hydrology* **256**: 120–141. DOI: 10.1016/S0022-1694(01)00529-7.
- Stoner AMK, Hayhoe K, Wuebbles DJ. 2009. Assessing general circulation model simulations of atmospheric teleconnection patterns. *Journal of Climate* **22**: 4348–4372. DOI: 10.1175/2009JCLI2577.1.
- Stott PA, Kettleborough JA. 2002. Origins and estimates of uncertainty in predictions of twenty-first century temperature rise. *Nature* **416**: 723–726. DOI: 10.1038/416723a.
- Sun Y, Solomon S, Dai A, Portman RW. 2006. How often does it rain? *Journal of Climate* **19**: 916–934. DOI: 10.1175/JCLI3672.1.
- USGCRP: National Assessment Synthesis Team. 2000. Climate change impacts on the United States: the potential consequences of climate variability and change. Cambridge University Press: Cambridge, UK. 620.
- USGCRP: Karl TR, Melillo JM, Peterson TC (eds). 2009. Global Climate Change Impacts in the United States. Cambridge University Press: Cambridge, UK; 196.
- Washington WM, Weatherly JW, Meehl GA, Semtner AJ, Bettge TW, Craig AP, Strand WG, Arblaster J, Wayland VB, James R, Zhang Y. 2000. Parallel climate model (PCM) control and transient simulations. *Climate Dynamics* **16**: 755–774. DOI: 10.1007/s003820000079.
- Widmann M, Bretherton CS, Salathe EP Jr. 2003. Statistical precipitation downscaling over the northwestern United States using numerically simulated precipitation as a predictor. *Journal of Climate* **16**: 799–816. DOI: 10.1175/1520-0442(2003)016<0799:SPDOTN>2.0.CO;2.
- Wilby RL, Wigley TML. 2000. Precipitation predictors for downscaling: observed and general circulation model relationships. *International Journal of Climatology* **20**: 641–661. DOI: 10.1002/(SICI)1097-0088(200005)20:6<641::AID-JOC501>3.0.CO;2-1.
- Wilby RL, Wigley TML, Conway D, Jones PD, Hewitson BC, Main J, Wilks DS. 1998. Statistical downscaling of general circulation model output: a comparison of methods. *Water Resources Research* **34**: 2995–3008. DOI: 10.1029/98WR02577.
- Wood AW, Leung LR, Sridhar V, Lettenmaier DP. 2004. Hydrologic implications of dynamical and statistical approaches to downscaling climate model outputs. *Climatic Change* **62**: 189–216. DOI: 10.1023/B:CLIM.0000013685.99609.9e.